

IBM MQ



Programming

Version 9 Release 0

Note

Before using this information and the product it supports, read the information in "Notices" on page 1321.

This edition applies to version 9 release 0 of IBM 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 v

Tables ix

Developing applications for IBM MQ . . . 1

Application development concepts	2
Actions that your applications can perform	3
Application programs using the MQI	5
Object-oriented applications	6
IBM MQ messages	9
Preparing and running Microsoft Transaction Server applications	41
Using IBM MQ with WebSphere Application Server	42
Design considerations for IBM MQ applications	43
Choosing to use IBM MQ classes for Java or IBM MQ classes for JMS.	46
Design techniques for messages	46
Selectors and message properties	47
Application design and performance considerations	48
Design techniques for advanced applications	50
Design and performance considerations for IBM i applications	51
Linux on POWER Systems - Little Endian applications	53
Design and performance considerations for z/OS applications	53
IMS and IMS bridge applications on IBM MQ for z/OS	57
Developing MQI applications with IBM MQ	70
IBM MQ data definition files	71
Writing a procedural application for queuing	74
Writing client procedural applications	285
User exits, API exits, and IBM MQ installable services	311
Building a procedural application	382
Handling procedural program errors	436
Multicast programming	442
Coding in C.	448
Coding in Visual Basic	450
Coding in COBOL.	451
Coding in System/390 assembler language (Message queue interface)	452
Coding IBM MQ programs in RPG (IBM i only)	455
Coding in pTAL	455
Coding in PL/I (z/OS only)	456
Using the IBM MQ sample procedural programs	457
Developing .NET applications	630
Getting started with IBM MQ classes for .NET	631
Writing and deploying IBM MQ .NET programs	647
Using the Component Object Model Interface (IBM MQ Automation Classes for ActiveX)	681

Designing and programming using IBM MQ Automation Classes for ActiveX	682
IBM MQ Automation Classes for ActiveX reference	688
Tracing IBM MQ Automation Classes for ActiveX	776
ActiveX interface to the MQAI	782
About the IBM MQ Automation Classes for ActiveX Starter samples	791
Developing C++ applications	796
C++ sample programs	798
C++ language considerations	803
Messaging in C++.	807
Building IBM MQ C++ programs.	814
Developing JMS and Java applications	824
Using IBM MQ classes for JMS	824
Using IBM MQ classes for Java	1021
Using the IBM MQ resource adapter	1092
Using the IBM MQ Headers package	1143
Setting up IBM MQ on IBM i with Java and JMS	1146
Developing applications for MQ Telemetry	1154
IBM MQ Telemetry Transport sample programs	1154
MQTT client programming concepts	1155
C client programming concepts	1168
Developing MQ Light applications	1168
AMQP and IBM MQ application interaction	1169
Developing Microsoft Windows Communication Foundation (WCF) applications with IBM MQ	1178
Introduction to the use of IBM MQ custom channel for WCF with .NET 3	1178
Using IBM MQ custom channels for WCF	1184
Using the WCF samples	1204
Problem determination on the WCF custom channel for IBM MQ	1210
Developing web services with IBM MQ	1217
Developing web services with IBM MQ transport for SOAP	1218
Developing REST applications with IBM MQ	1286
Messaging using the REST API	1288
Developing web services with IBM MQ bridge for HTTP	1293

Index 1303

Notices 1321

Programming interface information	1322
Trademarks	1323

Sending your comments to IBM 1325

Figures

1. Queues, messages, and applications	5	47. Complex deployment of intercepting subscribers	216
2. Representation of a message	9	48. Flow of application and trigger messages	236
3. Selection using MQSUB call	24	49. Relationship of queues within triggering	238
4. Selection using MQOPEN call	25	50. Queue manager groups	298
5. Standard Message Driven application consuming from two queues	35	51. MQCONN example	303
6. Single Threaded Message Driven application consuming from two queues	36	52. Understanding services, components, and entry points	316
7. Group of logical messages	36	53. UNIX and Linux authorization service stanzas in qm.ini	324
8. Segmented messages	37	54. Name service stanzas in qm.ini (for UNIX and Linux systems)	325
9. How distribution lists work.	123	55. Authorization service stanzas in qm.ini on IBM i	331
10. Opening a distribution list in C	125	56. Compile and link amqsaxe0.c on 32 bit Windows	338
11. Opening a distribution list in COBOL	125	57. Compile and link amqsaxe0.c on 64 bit Windows	338
12. Putting a message to a distribution list in C	127	58. Security exit loop	342
13. Putting a message to a distribution list in COBOL	127	59. Example of a send exit at the sender end of message channel	343
14. Logical order on a queue	135	60. Example of a receive exit at the receiver end of message channel	343
15. Physical order on a queue	136	61. Sample source code for a channel exit	348
16. Skipping backout using MQGMO_MARK_SKIP_BACKOUT	163	62. Sample DEF file for Windows	349
17. Simple IBM MQ publisher to a fixed topic.	175	63. Sender-initiated exchange with agreement	351
18. Sample output from first publisher example	176	64. Sender-initiated exchange with no agreement	352
19. Simple IBM MQ publisher to a variable topic.	179	65. Receiver-initiated exchange with agreement	353
20. Sample output from second publisher example	180	66. Receiver-initiated exchange with no agreement	354
21. Topic object associations	181	67. Security exit skeleton code	355
22. MQ publication consumer.	184	68. Fragments of JCL to link-edit the object module in the batch environment, using single-phase commit	422
23. Output from MQ publication consumer	185	69. Fragments of JCL to link-edit the object module in the batch environment, using two-phase commit	422
24. Managed MQ subscriber - part 1: declarations and parameter handling.	187	70. Fragments of JCL to link-edit the object module in the CICS environment	425
25. Managed MQ subscriber - part 2: code body.	189	71. Fragments of JCL to link-edit the object module in the IMS environment	426
26. MQ subscriber	190	72. Dynamic linking using COBOL in the batch environment	429
27. Unmanaged MQ subscriber - part 1: declarations.	195	73. Dynamic linking using COBOL in the CICS environment	429
28. Unmanaged MQ subscriber - part 2: parameter handling.	196	74. Dynamic linking using COBOL in the IMS environment	430
29. Unmanaged MQ subscriber - part 3: code body.	198	75. Dynamic linking using assembly language in the batch environment	430
30. Publish 130 to NYSE/IBM/PRICE	199	76. Dynamic linking using assembly language in the CICS environment	430
31. Receive the retained publication	199	77. Dynamic linking using assembly language in the IMS environment	430
32. Resume subscription	200	78. Dynamic linking using C language in the batch environment	431
33. Receive retained publication with new unmanaged non durable subscription	200		
34. Overlapping subscriptions	200		
35. Subscription topics cannot be changed	201		
36. Managed non-durable subscriber lifelines	203		
37. Managed durable subscriber lifelines	204		
38. Unmanaged durable subscriber lifelines	205		
39. Sequence of intercepting subscribers	208		
40. Preprocessor directives	211		
41. Declarations	211		
42. Initializations	212		
43. Preparing to intercept publications	213		
44. Intercept publication and republish	214		
45. Completion	214		
46. Interception and Publish exit in a cluster	215		

79. Dynamic linking using C language in the CICS environment	431	119. Sending a String in a JMSBytesMessage	906
80. Dynamic linking using C language in the IMS environment	431	120. Receiving a String from a JMSBytesMessage	906
81. Dynamic linking using PL/I in the batch environment	431	121. Send a JMSBytesMessage using a DataOutputStream	907
82. Dynamic linking using PL/I in the IMS environment	432	122. Receive a JMSBytesMessage using a DataInputStream	907
83. The database coordination samples	483	123. Incorrect code page conversion	910
84. Reconnectable client samples	494	124. Writing bytes representing a string in a JMSStreamMessage using the destination character set	911
85. File system import	501	125. Send text message in the character set defined by the destination	912
86. Running the Reference Message samples	506	126. Send text message in ccid 37	912
87. Request and Inquire samples using triggering	515	127. Receive text message	912
88. Sample IBM i Client/Server (Echo) program flowchart	518	128. Sending a number using the destination encoding in a JMSStreamMessage	912
89. Example of ubbstxcn.cfg file for IBM MQ for Windows	532	129. Sending a number using the destination encoding in a JMSBytesMessage	913
90. Sample TUXEDO makefile for IBM MQ for Windows	533	130. Sending a String in a JMSBytesMessage	913
91. Example of ubbstxcn.cfg file for IBM MQ for Windows	535	131. Receiving a String from a JMSBytesMessage	913
92. Sample TUXEDO makefile for IBM MQ for Windows	536	132. Enable queue manager data conversion	914
93. How TUXEDO samples work together	537	133. Set target coded character set for queue manager conversion	914
94. Example of a report from the Print Message sample application	594	134. Code snippet from amqsget0.c	915
95. Programs and panels for the TSO versions of the Mail Manager	600	135. RECORD.h	917
96. Programs and panels for the CICS version of the Mail Manager	601	136. Modify amqsget0.c	918
97. Example of a panel showing a list of waiting messages	602	137. TryMyRecord	926
98. Example of a panel showing the contents of a message	603	138. RECORD	927
99. Immediate Inquiry panel for the Credit Check sample application	605	139. MyRecord	928
100. Programs and queues for the Credit Check sample application (COBOL programs only)	607	140. EndPoint	929
101. Initial screen for Message Handler sample	618	141. MyProducer	930
102. Message list screen for Message Handler sample	618	142. MyConsumer	930
103. Chosen message is displayed	620	143. Inconsistently coded MQMD and message data	968
104. Supplied IBM MQ constants for encoding	685	144. Send a message with an MQ message body.	968
105. IBM MQ C++ classes (item handling)	796	145. ServerSessionPool and ServerSession functionality	1009
106. IBM MQ C++ classes (queue management)	797	146. The initial page of the IVT program	1136
107. JMS objects and their relationships	873	147. Page showing the results of a successful IVT	1137
108. How messages are transformed between JMS and IBM MQ using the MQRFH2 header	879	148. Page showing the results of an IVT that failed	1138
109. How JMS messages are transformed to IBM MQ messages with no MQRFH2 header	893	149. Successful IVT output	1141
110. Enable queue manager data conversion	898	150. WCF architecture for the SOAP/JMS interface	1183
111. Code snippet from amqsget0.c	901	151. WCF architecture for the Non-SOAP/Non-JMS interface	1184
112. Sending a String in a JMSBytesMessage	901	152. SOAP envelope	1223
113. Receiving a String from a JMSBytesMessage	901	153. Overview of IBM MQ transport for SOAP	1224
114. Sending and receiving a JMSObjectMessage	905	154. Queues used by SOAP/IBM MQ (single queue manager)	1225
115. Send text message in the character set defined by the destination	905	155. Queues used by SOAP/IBM MQ (separate queue managers)	1226
116. Send text message in ccid 37	905	156. Inline service: StockQuoteDotNet.asmx	1235
117. Receive text message	905	157. Code-behind: Design SQDNNonInline.asmx	1235
118. Send data in JMSStreamMessage and JMSMapMessage	906	158. Code-behind: Implementation: SQDNNonInline.asmx.cs	1236
		159. SQAxis2Axis.java	1245
		160. SQAxis2DotNet.java	1246
		161. WsdIClient.java	1247
		162. Static client using Eclipse generated proxy	1247
		163. Static client using amqwdeployWMQService generated proxy	1248

164. Dynamic client using Eclipse generated proxy	1248	184. Build command for .NET Framework V2 service	1273
165. Dynamic client using amqwdeployWMQService generated proxy	1248	185. Service definition	1277
166. DII client (No proxy)	1249	186. Axis 1.4 Java client	1277
167. SQA2StaticClient.java	1252	187. Client configuration and output	1278
168. SQA2DynamicClient.java	1253	188. Output from running SQA2AsyncClient	1280
169. SQA2CallbackHandler.java	1254	189. runpojo.bat: Windows, using a classpath	1280
170. SQA2AsyncClient.java	1255	190. runpojo.sh: Linux, using a classpath.	1280
171. Console output from SQA2AsyncClient.java	1255	191. runaxis2.bat: Windows, using axis2.bat	1281
172. DynamicProxyClientSync.java	1257	192. runaxis2.sh: Linux, using axis2.sh	1281
173. DynamicProxyClientAsyncPolling.java	1258	193. URL from StockQuoteAxis.wsd1	1282
174. DynamicProxyClientAsyncCallback.java	1259	194. Static client proxy constructor.	1284
175. DispatchClientSync.java	1261	195. Client program	1284
176. DispatchClientAsyncPolling.java	1263	196. Configuration and output	1285
177. DispatchClientAsyncCallback.java	1265	197. IBM MQ bridge for HTTP	1293
178. HTTP StockQuoteClientDotNet program	1267	198. Example of an HTTP POST request to a queue.	1294
179. Modified StockQuoteClientDotNet program	1268	199. Example of an HTTP POST response	1294
180. SQVB2Axis	1269	200. Example of an HTTP DELETE request	1295
181. SQVB2DotNet	1269	201. Example of an HTTP DELETE response	1295
182. SQCS2Axis	1270	202. Example of an HTTP GET request	1295
183. SQCS2DotNet	1270	203. Example of an HTTP GET response	1295

Tables

1. Boolean operator outcome when logic is A AND B	31	39. Installable service components summary	315
2. Boolean operator outcome when logic is A OR B	31	40. Example of entry-points for an installable service	322
3. Boolean operator outcome when logic is NOT A	31	41. Authorization service components summary	326
4. Mapping IBM MQ messages to IMS transaction types	64	42. Libraries that are now in the client and server packages	335
5. Library files for non-threaded AIX applications	80	43. Channel exits available for each channel type	341
6. Library files for threaded AIX applications	80	44. Libraries that are now in the client and server packages	365
7. Library files for non-threaded HP-UX applications	81	45. Sample cluster workload exit program location for Multiplatforms	380
8. Library files for threaded HP-UX applications	81	46. Essential code for CICS application programs on AIX: XA initialization routine	386
9. Library files for non-threaded IBM i applications	82	47. OSS headers	388
10. Library files for threaded IBM i applications	82	48. OSS public executable and public import libraries	389
11. Library files for non-threaded Linux applications	82	49. Guardian headers	389
12. Library files for threaded Linux applications	83	50. Guardian public executable and public import libraries	390
13. Library files for Solaris applications	83	51.	390
14. Library files for Windows applications	84	52.	391
15. The MQ_CONNECT_TYPE environment variable	99	53.	392
16. Resolving queue names when using MQOPEN.	103	54. Essential code for CICS applications (HP-UX)	397
17. How queue attributes and options of the MQOPEN call affect access to queues	107	55. Example of CRTPGM in the nonthreaded environment	404
18. MQPUT options relating to messages in groups and segments of logical messages	138	56. Example of CRTPGM in the threaded environment	405
19. Outcome when MQPUT or MQCLOSE call is not consistent with group and segment information	140	57. Non-threaded environment	405
20. Using message and correlation identifiers	146	58. Threaded environment	405
21. Using the group identifier	146	59. Essential code for CICS applications (Solaris)	411
22. MQGET options and read ahead	149	60. Location of IBM MQ libraries	413
23. Point to point versus publish/subscribe IBM MQ program pattern.	173	61. Context initiators and their associated context acceptors	418
24. Point to point vs. subscribe IBM MQ program patterns.	182	62. Side-deck name required for each coordination semantic.	421
25. Errors from MQSUB with different queue handles and subscription combinations	194	63. Variants of dynamic link libraries	423
26. Publication properties.	207	64. Call names for dynamic linking	428
27. Subscription options for intercepting subscribers	209	65. CICS adapter trace entries	433
28. MQPUT values for republished messages	210	66. MQI concepts and how they relate to multicast	442
29. Prefixes for selector names	219	67. Multicast event code descriptions	446
30. Versions of XA Switch	232	68. IBM MQ on UNIX and Linux sample programs demonstrating use of the MQI (C and COBOL).	459
31. z/OS environmental features	260	69. IBM MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL.	461
32. When to use a shared-initiation queue	269	70. IBM MQ for Windows sample programs demonstrating use of the MQI (C and COBOL)	465
33. Mapping IBM MQ messages to IMS transaction types	277	71. IBM MQ for Windows sample programs demonstrating use of the MQI (Visual Basic)	467
34. Programming languages supported in client environments	291	72. Sample programs demonstrating use of the MQI (C and COBOL) on IBM i.	468
35. Client system libraries on UNIX	308	73. Where to find the samples for IBM MQ on UNIX and Linux systems	473
36. Client system libraries on Windows systems	308		
37. Sample programs for UNIX client systems	309		
38. Sample programs for Windows client systems	309		

74. Where to find the samples for IBM MQ for Windows	474	116. mqext property name, synonym, data type, and folder	882
75. Permissible values for the property parameter	480	117. mqps property name, synonym, data type, and folder	882
76. Categories of legacy publish/subscribe sample C programs	499	118. MQRFH2 folders and properties used by JMS	882
77. Fields to change to directory paths	531	119. Property data types	883
78. Fields to change to directory paths	534	120. JMS header fields mapping to MQMD fields	884
79. IBM MQ client channel directory string attributes	559	121. JMS properties mapping to MQMD fields	884
80. IBM MQ client channel directory integer attributes	559	122. JMS provider-specific properties mapping to MQMD fields	885
81. IBM MQ client channel boolean attribute	560	123. Outgoing message field mapping	886
82. IBM MQ client channel list attributes	560	124. Outgoing message JMS property mapping	886
83. Source for the distributed queuing exit samples	573	125. Outgoing list message JMS provider-specific property mapping	887
84. Source for the data conversion exit samples (assembler language only)	573	126. Incoming message JMS header field mapping	891
85. Batch Put and Get samples	577	127. Incoming message property mapping	891
86. Batch Browse sample	577	128. Incoming message provider-specific JMS property mapping	892
87. Batch Print Message sample (C language only)	578	129. Message types and conversion types	900
88. Publish/Subscribe samples	578	130. Message types and conversion types	907
89. Other samples	578	131. Message types and conversion types	909
90. TSO Mail Manager sample	579	132. Supported conversions from one data type to another	936
91. TSO Message Handler sample	580	133. Property names and valid values for use in queue and topic URIs.	947
92. CICS Put and Get samples	582	134. Property names and descriptions	963
93. CICS Queue Attributes sample.	582	135. Property names, values, and set methods	963
94. CICS Mail Manager sample (COBOL only)	582	136. Property names, descriptions, and types	965
95. CICS Credit Check sample	583	137. Property names and descriptions	969
96. CICS Asynchronous Consumption and Publish/Subscribe samples	583	138. Property names, values, and set methods	969
97. Source and JCL for the Credit Check IMS sample (C only).	584	139. CipherSpecs supported by IBM MQ and their equivalent CipherSuites	976
98. Permissible values for the property parameter	592	140. CipherSuites and their supported and unsupported CipherSpecs	977
99. Input examples	626	141. The IBM MQ exits directory	980
100. IBM MQ and Microsoft.NET mapping table	668	142. Possible CLIENTRECCECTOPTIONS property values.	985
101. Location of sample applications for implementing TLS in managed .NET.	674	143. Possible CLIENTRECCECTOPTIONS property	987
102. Properties defined in the mqtrace.config file	680	144. PUTASYNCALLOWED and DEFPRESP properties determining if messages are put asynchronously.	993
103. IBM MQ string constants and their corresponding character strings.	683	145. READAHEADALLOWED and DEFREADA properties determining if read ahead is used when receiving or browsing non-persistent messages outside of a transaction.	994
104. Location of sample programs	799	146. Considerations when you are choosing between shared subscriptions and cloned subscriptions	1000
105. C/C++ header files	803	147. IBM MQ classes for Java installation directories	1028
106. z/OS sample program files	822	148. Samples directories	1028
107. IBM MQ classes for JMS installation directories	830	149. JRE directories	1029
108. Samples directories	831	150. CLASSPATH setting to run IBM MQ classes for Java applications, including the IBM MQ classes for Java sample applications.	1029
109. CLASSPATH setting to compile and run IBM MQ classes for JMS applications, including the sample applications	833	151. The location of the IBM MQ classes for Java libraries for each platform	1031
110. The location of the IBM MQ classes for JMS libraries for each platform	835	152. Which stanza of the client configuration file contains which attribute	1034
111. Which stanza of the client configuration file contains which attribute	838	153. The directory for channel exit programs	1065
112. Scripts provided with IBM MQ classes for JMS.	861		
113. The JMS domain independent and domain specific interfaces	873		
114. Possible values for NameValueCCSID field	881		
115. mcd property name, synonym, data type, and folder	881		

154. CipherSpecs supported by IBM MQ and their equivalent CipherSuites	1084	170. Key classes required for using the custom channel	1181
155. CipherSuites and their supported and unsupported CipherSpecs	1085	171. New binding element fields	1193
156. The directory containing wmq.jmsra.rar for each platform	1099	172. New binding configuration fields	1193
157. Installation of the resource adapter in an application server	1100	173. New client endpoint fields	1194
158. Properties of the ResourceAdapter object that are associated with the connection pool .	1105	174. New service endpoint fields	1194
159. Properties of an ActivationSpec object that are used to create a JMS connection.	1107	175. Values of binding properties when setting administratively or programmatically	1198
160. Properties of an ActivationSpec object that are used to create a JMS connection consumer	1115	176. IBM MQ required configuration	1208
161. Properties of a ConnectionFactory object	1122	177. WCF trace configuration variables	1212
162. Properties that are common to a Queue object and a Topic object	1131	178. WCF trace enablement combinations.	1213
163. Properties that are specific to a Queue object	1134	179. IBM MQ transport for SOAP application environments	1219
164. Properties that are specific to a Topic object	1134	180. Windows installation directories	1230
165. Synchronization behavior of methods that result in requests to the server	1158	181. AIX installation directories.	1230
166. Header field mappings	1173	182. HP-UX, Solaris, and Linux (all platforms) installation directories	1230
167. Properties field mappings	1174	183. Top down EJB web service configuration	1238
168. Message descriptor for AMQP message	1175	184. WebSphere JAX-WS JMS Binding Configuration	1238
169.	1176	185. Additional JNDI parameters	1283
		186. IBM MQ bridge for HTTP verbs	1294
		187. Queue manager configuration	1297
		188. Set or modify the following fields	1298
		189. Publish/subscribe configuration modes	1299
		190. Location of HTTP samples.	1300





Developing applications for IBM MQ

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM® MQ support applications written in procedural languages, and object oriented languages and frameworks.

Before you develop applications for IBM MQ, ensure you are familiar with the concepts in IBM MQ Technical overview.

Support for procedural programming languages

IBM MQ provides support for the following procedural programming languages:

- C
- COBOL
-  Visual Basic (Windows systems only)
-  Assembler language (IBM MQ for z/OS® only)
-  RPG (IBM MQ for IBM i only)
-  PL/I (IBM MQ for z/OS only)

These languages use the message queue interface (MQI) to access message queuing services. For more information, see “Developing MQI applications with IBM MQ” on page 70.

Support for object oriented languages and frameworks



IBM MQ also provides support for the following object-oriented programming languages and frameworks:

- .NET
- ActiveX
- C++
- Java™
- JMS

These languages and frameworks use the IBM MQ Object Model, which provides classes that provide the same functionality as IBM MQ calls and structures, but that are a more natural way of programming in an object-oriented environment. Some of the languages and frameworks that use the IBM MQ Object Model provide additional functions that are not available to the procedural languages using the MQI. For more information, see “Object-oriented applications” on page 6.

Support for programmatic REST APIs

IBM MQ provides support for the following programmatic REST APIs to send and receive messages:

-  IBM MQ messaging REST API
-  IBM MQ z/OS Connect EE
- IBM Integration Bus
- IBM DataPower Gateway
- IBM MQ bridge for HTTP

Related information:

Securing

Application development concepts

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

Before you start to design and write your IBM MQ applications, familiarize yourself with the basic IBM MQ concepts, see the topics in Technical overview. For information about the types of application you can write for IBM MQ, see “Developing applications for IBM MQ” on page 1.

Use the following links to find out about IBM MQ concepts specific to application development:

Related concepts:

“Using the MQI in a client application” on page 286

This collection of topics considers the differences between writing your IBM MQ application to run in a message queue interface (MQI) client environment and to run in the full IBM MQ queue manager environment.

“Channel-exit programs for messaging channels” on page 341

This collection of topics contains information about IBM MQ channel-exit programs for messaging channels.

“Design considerations for IBM MQ applications” on page 43

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

“Writing a procedural application for queuing” on page 74

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

“Writing client procedural applications” on page 285

What you need to know to write client applications on IBM MQ using a procedural language.

“Developing MQI applications with IBM MQ” on page 70

IBM MQ provides support for C, Visual Basic, COBOL, Assembler, RPG, pTAL, and PL/I. These procedural languages use the message queue interface (MQI) to access message queuing services.

“Object-oriented applications” on page 6

IBM MQ provides support for .NET, ActiveX, C++, Java, and JMS. These languages and frameworks use the IBM MQ Object Model, which provides classes that provide the same functionality as IBM MQ calls and structures. Some of the languages and frameworks that use the IBM MQ Object Model provide additional functions that are not available when you use procedural languages with the message queue interface (MQI).

“Using IBM MQ classes for JMS” on page 824

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the javax.jms package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

“Using the Component Object Model Interface (IBM MQ Automation Classes for ActiveX)” on page 681

The IBM MQ Automation Classes for ActiveX (MQAX) are ActiveX components that provide classes that you can use in your application to access IBM MQ.

“Using IBM MQ classes for Java” on page 1021

Use IBM MQ in a Java environment. IBM MQ classes for Java allow a Java application to connect to IBM MQ as an IBM MQ client, or connect directly to an IBM MQ queue manager.

“Developing .NET applications” on page 630

IBM MQ classes for .NET allow a program written in the .NET programming framework to connect to IBM MQ as an IBM MQ MQI client or to connect directly to an IBM MQ server.

“Developing C++ applications” on page 796

IBM MQ provides C++ classes equivalent to IBM MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI.

“Building a procedural application” on page 382

You can write an IBM MQ application in one of several procedural languages, and run the application on several different platforms.

Related tasks:

“Developing web services with IBM MQ” on page 1217

You can develop IBM MQ applications for web services using the IBM MQ transport for SOAP.

“Using the IBM MQ sample procedural programs” on page 457

These sample programs are written in procedural languages, and demonstrate typical uses of the Message Queue Interface (MQI). IBM MQ programs on different platforms.

Related information:

Transactional support scenarios

Actions that your applications can perform

You can develop applications to send and receive messages that you need to support your business processes. You can also develop applications to manage your queue managers and related resources.

Multi

Actions that your applications can perform on IBM MQ for Multiplatforms

On Multiplatforms, you can write applications that perform the following actions:

- Send messages to other applications running under the same operating systems. The applications can be on either the same or another system.
- Send messages to applications that run on other IBM MQ platforms.
- Use message queuing from within CICS® for **IBM i** IBM i, TXSeries® for AIX®, HP-UX, Solaris, and Windows systems.
- Use message queuing from within Encina for AIX, HP-UX, Solaris, and Windows systems.
- Use message queuing from within Tuxedo for AIX, AT&T, HP-UX, Solaris, and Windows systems.
- Use IBM MQ as a transaction manager, coordinating updates made by external resource managers within IBM MQ units of work. The following external resource managers are supported and comply with the X/OPEN XA interface
 - DB2®
 - Informix®
 - Oracle
 - Sybase
- Process several messages together as a single unit of work that can be committed or backed out.
- Run from a full IBM MQ environment, or run from an IBM MQ client environment.

z/OS

Actions that your applications can perform on IBM MQ for z/OS

On z/OS, you can write applications that perform the following actions:


- Use message queuing within CICS or IMS™.
- Send messages between batch, CICS, and IMS applications, selecting the most appropriate environment for each function.
- Send messages to applications that run on other IBM MQ platforms.
- Process several messages together as a single unit of work that can be committed or backed out.

- Send messages to, and interact with, IMS applications by means of the IMS bridge.
- Participate in units of work coordinated by RRS.

Each environment within z/OS has its own characteristics, advantages, and disadvantages. The advantage of IBM MQ for z/OS is that applications are not tied to any one environment, but can be distributed to take advantage of the benefits of each environment. For example, you can develop end-user interfaces using TSO or CICS, you can run processing-intensive modules in z/OS batch, and you can run database applications in IMS or CICS. In all cases, the various parts of the application can communicate using messages and queues.

Designers of IBM MQ applications must be aware of the differences and limitations imposed by these environments. For example:

- IBM MQ provides facilities that allow intercommunication between queue managers (this is known as *distributed queuing*).
- Methods of committing and backing out changes differ between the batch and CICS environments.
- IBM MQ for z/OS provides support in the IMS environment for online message processing programs (MPPs), interactive fast path programs (IFPs), and batch message processing programs (BMPs). If you are writing batch DL/I programs, follow the guidance given in topics such as “Building z/OS batch applications” on page 421 and “z/OS batch considerations” on page 87 for z/OS batch programs.
- Although multiple instances of IBM MQ for z/OS can exist on a single z/OS system, a CICS region can connect to only one queue manager at a time. However, more than one CICS region can be connected to the same queue manager. In the IMS and z/OS batch environments, programs can connect to more than one queue manager.
- IBM MQ for z/OS allows local queues to be shared by a group of queue managers, giving improved throughput and availability. Such queues are called *shared queues*, and the queue managers form a *queue-sharing group*, which can process messages on the same shared queues. Batch applications can connect to one of several queue managers within a queue-sharing group by specifying the queue-sharing group name, instead of a particular queue manager name. This is known as *group batch attach*, or more simply *group attach*. See Shared queues and queue-sharing groups.

 The differences between the supported environments, and their limitations, are explained further in “Using and writing applications on IBM MQ for z/OS” on page 258.

Related concepts:

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

“Design considerations for IBM MQ applications” on page 43

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

“Writing a procedural application for queuing” on page 74

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

“Writing client procedural applications” on page 285

What you need to know to write client applications on IBM MQ using a procedural language.

“Using IBM MQ classes for JMS” on page 824

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the `javax.jms` package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

“Using the Component Object Model Interface (IBM MQ Automation Classes for ActiveX)” on page 681

The IBM MQ Automation Classes for ActiveX (MQAX) are ActiveX components that provide classes that you can use in your application to access IBM MQ.

“Using IBM MQ classes for Java” on page 1021

Use IBM MQ in a Java environment. IBM MQ classes for Java allow a Java application to connect to IBM MQ as an IBM MQ client, or connect directly to an IBM MQ queue manager.

“Developing .NET applications” on page 630

IBM MQ classes for .NET allow a program written in the .NET programming framework to connect to IBM MQ as an IBM MQ MQI client or to connect directly to an IBM MQ server.

“Developing Microsoft Windows Communication Foundation (WCF) applications with IBM MQ” on page 1178

The Microsoft Windows Communication Foundation (WCF) custom channel for IBM MQ sends and receives messages between WCF clients and services.

“Developing C++ applications” on page 796

IBM MQ provides C++ classes equivalent to IBM MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI.

“Building a procedural application” on page 382

You can write an IBM MQ application in one of several procedural languages, and run the application on several different platforms.

Related tasks:

“Using the IBM MQ sample procedural programs” on page 457

These sample programs are written in procedural languages, and demonstrate typical uses of the Message Queue Interface (MQI). IBM MQ programs on different platforms.

Related information:

Securing

Application programs using the MQI

IBM MQ application programs need certain objects before they can run successfully.

Figure 1 shows an application that removes messages from a queue, processes them, and then sends some results to another queue on the same queue manager.




Figure 1. Queues, messages, and applications

Whereas applications can put messages onto local or remote queues (using MQPUT), they can only get messages directly from local queues (using MQGET).

Before this application can run, the following conditions must be satisfied:

- The queue manager must exist and be running.
- The first application queue, from which the messages are to be removed, must be defined.
- The second queue, on which the application puts the messages, must also be defined.

- The application must be able to connect to the queue manager. To do this it must be linked to IBM MQ. See “Building a procedural application” on page 382.
- The applications that put the messages on the first queue must also connect to a queue manager. If they are remote, they must also be set up with transmission queues and channels. This part of the system is not shown in Figure 1 on page 5.

Object-oriented applications


IBM MQ provides support for .NET, ActiveX, C++, Java, and JMS. These languages and frameworks use the IBM MQ Object Model, which provides classes that provide the same functionality as IBM MQ calls and structures. Some of the languages and frameworks that use the IBM MQ Object Model provide additional functions that are not available when you use procedural languages with the message queue interface (MQI).

For details of the classes, methods and properties provided by this model, see “The IBM MQ Object Model” on page 7.

.NET See Developing .NET applications for information about coding .NET programs using the IBM MQ .NET classes. Message Service Clients for C/C++ and .NET provide an application programming Interface (API) called XMS that has the same set of interfaces as the Java Message Service (JMS) API.

ActiveX

The IBM MQ ActiveX is commonly known as the MQAX. The MQAX is included as part of IBM MQ for Windows. Support for ActiveX has been stabilized at the IBM WebSphere® MQ Version 6.0 level. For information about coding programs using the IBM MQ Object Model in ActiveX Using the Component Object Model Interface (WebSphere MQ Automation Classes for ActiveX).

 From Version 9.0, IBM MQ support for MicrosoftActive X is deprecated. IBM MQ classes for .NET are the recommended replacement technology. For more information, see Developing .NET applications.

C++ IBM MQ provides C++ classes equivalent to IBM MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI. See Using C++ for information about coding programs using the IBM MQ Object Model in C++. Message Service Clients for C/C++ and .NET provide an application programming Interface (API) called XMS that has the same set of interfaces as the Java Message Service (JMS) API.

Java See Using IBM MQ classes for Java for information about coding programs using the IBM MQ Object Model in Java. IBM will make no further enhancements to the IBM MQ classes for Java and they are functionally stabilized at the level shipped in IBM MQ Version 8.0. For information about the differences between the IBM MQ classes for Java and the IBM MQ classes for JMS to help you decide which to use, see “Choosing to use IBM MQ classes for Java or IBM MQ classes for JMS” on page 46.

JMS IBM MQ also provides classes that implement the Java Message Service (JMS) specification. For details of the IBM MQ classes for JMS, see Using IBM MQ classes for JMS. For information about the differences between IBM MQ classes for Java and IBM MQ classes for JMS to help you decide which to use, see “Choosing to use IBM MQ classes for Java or IBM MQ classes for JMS” on page 46.

IBM Message Service Client for C/C++ and IBM Message Service Client for .NET provide an application programming interface (API) called XMS that has the same set of interfaces as the Java Message Service (JMS) API. For more information, see Introduction to the IBM Message Service Client for .NET.

Related concepts:

“Developing MQI applications with IBM MQ” on page 70

IBM MQ provides support for C, Visual Basic, COBOL, Assembler, RPG, pTAL, and PL/I. These procedural languages use the message queue interface (MQI) to access message queuing services.

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

Related information:

Technical overview

Application development reference

The IBM MQ Object Model

The IBM MQ Object Model consists of classes, methods and properties.

The IBM MQ Object Model consists of:

- *Classes* representing familiar IBM MQ concepts such as queue managers, queues, and messages.
- *Methods* on each class corresponding to MQI calls.
- *Properties* on each class corresponding to attributes of IBM MQ objects.

When creating an IBM MQ application using the IBM MQ Object Model, you create instances of these classes in the application. An instance of a class in object-oriented programming is called an *object*. When an object has been created, you interact with the object by examining or setting the values of the object's properties (the equivalent of issuing an MQINQ or MQSET call), and by making method calls against the object (the equivalent of issuing the other MQI calls).

Classes

The IBM MQ Object Model provides the following base set of classes.

The actual implementation of the model varies slightly between the different supported object-oriented environments.

MQQueueManager

An object of the MQQueueManager class represents a connection to a queue manager. It has methods to Connect(), Disconnect(), Commit(), and Backout() (the equivalent of MQCONN or MQCONNX, MQDISC, MQCMIT, and MQBACK). It has properties corresponding to the attributes of a queue manager. Accessing a queue manager attribute property implicitly connects to the queue manager if not already connected. Destroying an MQQueueManager object implicitly disconnects from the queue manager.

MQQueue

An object of the MQQueue class represents a queue. It has methods to Put() and Get() messages to and from the queue (the equivalent of MQPUT and MQGET). It has properties corresponding to the attributes of a queue. Accessing a queue attribute property, or issuing a Put() or Get() method call, implicitly opens the queue (the equivalent of MQOPEN). Destroying an MQQueue object implicitly closes the queue (the equivalent of MQCLOSE).

MQTopic

An object of the MQTopic class represents a topic. It has methods to Put() (publish) and Get() (receive or subscribe) messages to and from the topic (the equivalent of MQPUT and MQGET). It has properties corresponding to the attributes of a topic. An MQTopic object can only be accessed for publication or subscription, not both simultaneously. When used for receiving messages the MQTopic object can be created with an unmanaged or managed subscription and as a durable or non-durable subscriber - multiple overloaded constructors are provided for these differing scenarios.

MQMessage

An object of the MQMessage class represents a message to be put on a queue or got from a queue. It contains a buffer, and encapsulates both application data and MQMD. It has properties corresponding to MQMD fields, and methods that allow you to write and read user data of different types (for example, strings, long integers, short integers, single bytes) to and from the buffer.

MQPutMessageOptions

An object of the MQPutMessageOptions class represents the MQPMO structure. It has properties corresponding to MQPMO fields.

MQGetMessageOptions

An object of the MQGetMessageOptions class represents the MQGMO structure. It has properties corresponding to MQGMO fields.

MQProcess

An object of the MQProcess class represents a process definition (used with triggering). It has properties that represent the attributes of a process definition.

Multi

MQDistributionList

An object of the MQDistributionList class represents a distribution list (used to send multiple messages with a single MQPUT). It contains a list of MQDistributionListItem objects.

Multi

MQDistributionListItem

An object of the MQDistributionListItem class represents a single distribution list destination. It encapsulates the MQOR, MQRR, and MQPMR structures, and has properties corresponding to the fields of these structures.

Object references

In an IBM MQ program that uses the MQI, IBM MQ returns connection handles and object handles to the program.

These handles must be passed as parameters on subsequent IBM MQ calls. With the IBM MQ Object Model, these handles are hidden from the application program. Instead, the creation of an object from a class results in an object reference being returned to the application program. It is this object reference that is used when making method calls and property accesses against the object.

Return codes

Issuing a method call or setting a property value results in return codes being set.

These return codes are a completion code and a reason code, and are themselves properties of the object. The values of completion code and reason code are the same as those defined for the MQI, with some extra values specific to the object-oriented environment.

IBM MQ messages

An IBM MQ message consists of message properties and application data. The message queuing message descriptor (MQMD) contains the control information that accompanies the application data when a message travels between the sending and receiving applications.

Parts of a message

IBM MQ messages consist of two parts:

- Message properties
- Application data

Figure 2 represents a message and shows how it is logically divided into message properties and application data.




Figure 2. Representation of a message

The application data that is carried in an IBM MQ message is not changed by a queue manager unless data conversion is carried out on it. Also, IBM MQ does not put any restrictions on the content of this data. The length of the data in each message cannot exceed the value of the **MaxMsgLength** attribute of both the queue and queue manager.

ULW On UNIX, Linux, and Windows, the *MaxMsgLength* defaults to 100 MB (104 857 600 bytes).

IBM i On IBM i, the *MaxMsgLength* defaults to 100 MB (104 857 600 bytes). If you are intending to use IBM MQ messages greater than 15 MB on IBM i, see “Building your procedural application on IBM i” on page 404.

z/OS On z/OS, the **MaxMsgLength** attribute of the queue manager is fixed at 100 MB and the **MaxMsgLength** attribute of the queue defaults to 4 MB (4 194 304 bytes) which you can change up to a maximum of 100 MB if required.

Make your messages slightly shorter than the value of the **MaxMsgLength** attribute in some circumstances. For more information, see “The data in your message” on page 116.

You create a message when you use the MQPUT or MQPUT1 MQI calls. As input to these calls, you supply the control information (such as the priority of the message and the name of a reply queue) and your data, and the call then puts the message on a queue. See MQPUT and MQPUT1 for more information about these calls.

Message descriptor

You can access message control information by using the MQMD structure, which defines the *message descriptor*.

For a full description of the MQMD structure, see MQMD - Message descriptor.

For a description of how to use the fields within the MQMD that contain information about the origin of the message, see "Message context" on page 40.

There are different versions of the message descriptor. Additional information for grouping and segmenting messages (see "Message groups" on page 36) is provided in Version 2 of the message descriptor (or the MQMDE). This is the same as the Version 1 message descriptor but has extra fields. These fields are described in the MQMDE - Message descriptor extension.

Types of message

There are four types of messages defined by IBM MQ.

These four messages are:

- Datagram
- Request messages
- Reply messages
- Report messages
 - Types of report message
 - Report message options

Applications can use the first three types of messages to pass information between themselves. The fourth type, report, is for applications and queue managers to use to report information about events such as the occurrence of an error.

Each type of message is identified by an MQMT_* value. You can also define your own types of message. For the range of values you can use, see MsgType.

Datagrams

Use a *datagram* when you do not require a reply from the application that receives the message (that is, gets the message from the queue).

An example of an application that might use datagrams is one that displays flight information in an airport lounge. A message might contain the data for a whole screen of flight information. Such an application is unlikely to request an acknowledgment for a message because it probably does not matter if a message is not delivered. The application sends an update message after a short time.

Request messages

Use a *request message* when you want a reply from the application that receives the message.

An example of an application that could use request messages is one that displays the balance of a checking account. The request message could contain the number of the account, and the reply message would contain the account balance.

If you want to link your reply message with your request message, there are two options:

- Make the application that handles the request message responsible for ensuring that it puts information into the reply message that relates to the request message.

- Use the report field in the message descriptor of your request message to specify the content of the *MsgId* and *CorrelId* fields of the reply message:
 - You can request that either the *MsgId* or the *CorrelId* of the original message is to be copied into the *CorrelId* field of the reply message (the default action is to copy *MsgId*).
 - You can request that either a new *MsgId* is generated for the reply message, or that the *MsgId* of the original message is to be copied into the *MsgId* field of the reply message (the default action is to generate a new message identifier).

Reply messages

Use a *reply message* when you reply to another message.

When you create a reply message, respect any options that were set in the message descriptor of the message to which you are replying. Report options specify the content of the message identifier (*MsgId*) and correlation identifier (*CorrelId*) fields. These fields allow the application that receives the reply to correlate the reply with its original request.

Report messages

Report messages inform applications about events such as the occurrence of an error when processing a message.

They can be generated by:

- A queue manager,
- A message channel agent (for example, if they cannot deliver the message), or
- An application (for example, if it cannot use the data in the message).

Report messages can be generated at any time, and might arrive on a queue when your application is not expecting them.

Types of report message

When you put a message on a queue, you can select to receive:

- An *exception report message*. This is sent in response to a message with the exceptions flag set. It is generated by the message channel agent (MCA) or the application.
- An *expiry report message*. This indicates that an application attempted to retrieve a message that had reached its expiry threshold; the message is marked to be discarded. This type of report is generated by the queue manager.
- A *confirmation of arrival (COA) report message*. This indicates that the message has reached its target queue. It is generated by the queue manager.
- A *confirmation of delivery (COD) report message*. This indicates that the message has been retrieved by a receiving application. It is generated by the queue manager.
- A *positive action notification (PAN) report message*. This indicates that a request has been successfully serviced (that is, the action requested in the message has been performed successfully). This type of report is generated by the application.
- A *negative action notification (NAN) report message*. This indicates that a request has not been successfully serviced (that is, the action requested in the message has not been performed successfully). This type of report is generated by the application.

Note: Each type of report message contains one of the following:

- The entire original message
- The first 100 bytes of data in the original message
- No data from the original message

You can request more than one type of report message when you put a message on a queue. If you select the delivery confirmation report message and the exception report message options, if the message fails to be delivered, you receive an exception report message. However, if you select only the delivery confirmation report message option and the message fails to be delivered, you do not get an exception report message.

The report messages that you request, when the criteria for generating a particular message are met, are the only ones that you receive.

Report message options

You can *discard* a message after an exception has arisen. If you select the discard option, and have requested an exception report message, the report message goes to the *ReplyToQ* and *ReplyToQMgr*, and the original message is discarded.

Note: A benefit of this is that you can reduce the number of messages going to the dead-letter queue. However, it does mean that your application, unless it sends only datagram messages, has to deal with returned messages. When an exception report message is generated, it inherits the persistence of the original message.

If a report message cannot be delivered (if the queue is full, for instance), the report message is placed on the dead-letter queue.

If you want to receive a report message, specify the name of your reply-to queue in the *ReplyToQ* field; otherwise the MQPUT or MQPUT1 of your original message fails with MQRC_MISSING_REPLY_TO_Q.

You can use other report options in the message descriptor (MQMD) of a message to specify the content of the *MsgId* and *CorrelId* fields of any report messages that are created for the message:

- You can request that either the *MsgId* or the *CorrelId* of the original message is to be copied into the *CorrelId* field of the report message. The default action is to copy the message identifier. Use MQRO_COPY_MSG_ID_TO_CORRELID because it enables the sender of a message to correlate the reply or report message with the original message. The correlation identifier of the reply or report message is identical to the message identifier of the original message.
- You can request that either a new *MsgId* is generated for the report message, or that the *MsgId* of the original message is to be copied into the *MsgId* field of the report message. The default action is to generate a new message identifier. Use MQRO_NEW_MSG_ID because it ensures that each message in the system has a different message identifier, and can be distinguished unambiguously from all other messages in the system.
- Specialized applications might need to use MQRO_PASS_MSG_ID or MQRO_PASS_CORREL_ID. However, you need to design the application that reads the messages from the queue to ensure that it works correctly when, for example, the queue contains multiple messages with the same message identifier.

Server applications must check the settings of these flags in the request message, and set the *MsgId* and *CorrelId* fields in the reply or report message appropriately.

Applications that act as intermediaries between a requester application and a server application do not need to check the settings of these flags. This is because these applications typically need to forward the message to the server application with the *MsgId*, *CorrelId*, and *Report* fields unchanged. This allows the server application to copy the *MsgId* from the original message in the *CorrelId* field of the reply message.

When generating a report about a message, server applications must test to see if any of these options have been set.

For more information about how to use report messages, see Report.

To indicate the nature of the report, queue managers use a range of feedback codes. They put these codes in the *Feedback* field of the message descriptor of a report message. Queue managers can also return MQI reason codes in the *Feedback* field. IBM MQ defines a range of feedback codes for applications to use.

For more information about feedback and reason codes, see [Feedback](#).

An example of a program that could use a feedback code is one that monitors the workloads of other programs serving a queue. If there is more than one instance of a program serving a queue, and the number of messages arriving on the queue no longer justifies this, such a program can send a report message (with the feedback code MQFB_QUIT) to one of the serving programs to indicate that the program should terminate its activity. (A monitoring program could use the MQINQ call to find out how many programs are serving a queue.)

Reports and segmented messages:

Not supported on IBM MQ for z/OS.

If a message is segmented and you ask for reports to be generated, you might receive more reports than you would have done had the message not been segmented.

For a description of segmented messages, see “Message segmentation” on page 153.

For reports generated by IBM MQ

If you segment your messages or allow the queue manager to do so, there is only one case in which you can expect to receive a single report for the entire message. This is when you have requested only COD reports, and you have specified MQGMO_COMPLETE_MSG on the getting application.

In other cases your application must be prepared to deal with several reports; usually one for each segment.

Note: If you segment your messages, and you need only the first 100 bytes of the original message data to be returned, change the setting of the report options to ask for reports with no data for segments that have an offset of 100 or more. If you do not do this, and you leave the setting so that each segment requests 100 bytes of data, and you retrieve the report messages with a single MQGET specifying MQGMO_COMPLETE_MSG, the reports assemble into a large message containing 100 bytes of read data at each appropriate offset. If this happens, you need a large buffer or you need to specify MQGMO_ACCEPT_TRUNCATED_MSG.

For reports generated by applications

If your application generates reports, always copy the IBM MQ headers that are present at the start of the original message data to the report message data.

Then add none, 100 bytes, or all of the original message data (or whatever other amount you would usually include) to the report message data.

You can recognize the IBM MQ headers that must be copied by looking at the successive Format names, starting with the MQMD and continuing through any headers present. The following Format names indicate these IBM MQ headers:

- MQMDE
- MQDLH
- MQXQH
- MQIIH

- MQH*

MQH* means any name that starts with the characters MQH.

The Format name occurs at specific positions for MQDLH and MQXQH, but for the other IBM MQ headers it occurs at the same position. The length of the header is contained in a field that also occurs at the same position for MQMDE, MQIMS, and all MQH* headers.

If you are using a Version 1 MQMD, and you are reporting on a segment, or a message in a group, or a message for which segmentation is allowed, the report data must start with an MQMDE. Set the *OriginalLength* field to the length of the original message data excluding the lengths of any IBM MQ headers that you find.

Retrieving reports

If you ask for COA or COD reports, you can ask for them to be reassembled for you with MQGMO_COMPLETE_MSG.

An MQGET with MQGMO_COMPLETE_MSG is satisfied when enough report messages (of a single type, for example COA, and with the same *GroupId*) are present on the queue to represent one complete original message. This is true even if the report messages themselves do not contain the complete original data; the *OriginalLength* field in each report message gives the length of original data represented by that report message, even if the data itself is not present.

You can use this technique even if there are several different report types present on the queue (for example, both COA and COD), because an MQGET with MQGMO_COMPLETE_MSG reassembles report messages only if they have the same *Feedback* code. However, you cannot usually use this technique for exception reports, because, in general, these have different *Feedback* codes.

You can use this technique to get a positive indication that the entire message has arrived. However, in most circumstances you need to cater for the possibility that some segments arrive while others might generate an exception (or expiry, if you have allowed this). You cannot use MQGMO_COMPLETE_MSG in this case, because, in general, you might get different *Feedback* codes for different segments and, you might get more than one report for a segment. You can, however, use MQGMO_ALL_SEGMENTS_AVAILABLE.

To allow for this you might need to retrieve reports as they arrive, and build up a picture in your application of what happened to the original message. You can use the *GroupId* field in the report message to correlate reports with the *GroupId* of the original message, and the *Feedback* field to identify the type of each report message. The way in which you do this depends on your application requirements.

One approach is as follows:

- Ask for COD reports and exception reports.
- After a specific time, check whether a complete set of COD reports has been received using MQGMO_COMPLETE_MSG. If so, your application knows that the entire message has been processed.
- If not, and exception reports relating to this message are present, handle the problem as for unsegmented messages, but ensure that you clean up orphan segments at some point.
- If there are segments for which there are no reports of any kind, the original segments (or the reports) might be waiting for a channel to be reconnected, or the network might be overloaded at some point. If no exception reports at all have been received (or if you think that the ones you have might be temporary only), you might decide to let your application wait a little longer.

As before, this is similar to the considerations you have when dealing with unsegmented messages, except that you must also consider the possibility of cleaning up orphan segments.

If the original message is not critical (for example, if it is a query, or a message that can be repeated later), set an expiry time to ensure that orphan segments are removed.

Back-level queue managers

When a report is generated by a queue manager that supports segmentation, but is received on a queue manager that does not support segmentation, the MQMDE structure (which identifies the *Offset* and *OriginalLength* represented by the report) is always included in the report data, in addition to zero, 100 bytes, or all of the original data in the message.

However, if a segment of a message passes through a queue manager that does not support segmentation, if a report is generated there, the MQMDE structure in the original message is treated purely as data. It is not therefore included in the report data if zero bytes of the original data have been requested. Without the MQMDE, the report message might not be useful.

Request at least 100 bytes of data in reports if there is a possibility that the message might travel through a back-level queue manager.

Format of message control information and message data

The queue manager is only interested in the format of the control information within a message, whereas applications that handle the message are interested in the format of both the control information and the data.

Format of message control information

Control information in the character-string fields of the message descriptor must be in the character set used by the queue manager.

The **CodedCharSetId** attribute of the queue manager object defines this character set. Control information must be in this character set because, when applications pass messages from one queue manager to another, message channel agents that transmit the messages use the value of this attribute to determine what data conversion to perform.

Format of message data

You can specify any of the following things:

- The format of the application data
- The character set of the character data
- The format of numeric data

To do this, use these fields:

Format This indicates to the receiver of a message the format of the application data in the message.

When the queue manager creates a message, in some circumstances it uses the *Format* field to identify the format of that message. For example, when a queue manager cannot deliver a message, it puts the message on a dead-letter (undelivered message) queue. It adds a header (containing more control information) to the message, and changes the *Format* field to show this.

The queue manager has a number of *built-in formats* with names beginning MQ, for example MQFMT_STRING. If these do not meet your needs, you can define your own formats (*user-defined formats*), but you must not use names beginning with MQ for these.

When you create and use your own formats, you must write a data-conversion exit to support a program getting the message using MQGMO_CONVERT.

CodedCharSetId

This defines the character set of character data in the message. If you want to set this character set to that of the queue manager, you can set this field to the constant MQCCSI_Q_MGR or MQCCSI_INHERIT.

When you get a message from a queue, compare the value of the *CodedCharSetId* field with the value that your application is expecting. If the two values differ, you might need to convert any character data in the message or use a data-conversion message exit if one is available.

Encoding

This describes the format of numeric message data that contains binary integers, packed-decimal integers, and floating point numbers. It is typically encoded according to the particular machine on which the queue manager is running.

When you put a message on a queue, you typically specify the constant MQENC_NATIVE in the *Encoding* field. This means that the encoding of your message data is the same as that of the machine on which your application is running.

When you get a message from a queue, compare the value of the *Encoding* field in the message descriptor with the value of the constant MQENC_NATIVE on your machine. If the two values differ, you might need to convert any numeric data in the message or use a data-conversion message exit if one is available.

Application data conversion:

Application data might need to be converted to the character set and the encoding required by another application where different platforms are concerned.

It can be converted at the sending queue manager, or at the receiving queue manager. If the library of built-in formats does not meet your needs, you can define your own. The type of conversion depends on the message format that is specified in the format field of the message descriptor, MQMD.

Note: Messages with MQFMT_NONE specified are not converted.

Conversion at the sending queue manager

Set the CONVERT channel attribute to YES if you need the sending message channel agent (MCA) to convert the application data.

The conversion is performed at the sending queue manager for certain built-in formats and for user-defined formats if a suitable user exit is supplied.

Built-in formats

These include:

- Messages that are all characters (using the format name MQFMT_STRING)
- IBM MQ defined messages, for example Programmable Command Formats

IBM MQ uses Programmable Command Format messages for administration messages and events (the format name used is MQFMT_ADMIN in this case). You can use the same format (using the format name MQFMT_PCF) for your own messages, and take advantage of the built-in data conversion.

The queue manager built-in formats all have names beginning with MQFMT. They are listed and described in Format.

Application-defined formats

For user-defined formats, application data conversion must be performed by a data-conversion exit program (for more information, see "Writing data-conversion exits" on page 364). In a client-server environment, the exit is loaded at the server and conversion takes place there.

Conversion at the receiving queue manager

Application message data can be converted by the receiving queue manager for both built-in and user-defined formats.

The conversion is performed during the processing of an MQGET call if you specify the MQGMO_CONVERT option. For details, see the Options

Coded character sets

IBM MQ products support the coded character sets that are provided by the underlying operating system.

When you create a queue manager, the queue manager coded character set ID (CCSID) used is based on that of the underlying environment. If this is a mixed code page, IBM MQ uses the SBCS part of the mixed code page as the queue manager CCSID.

For general data conversion, if the underlying operating system supports DBCS code pages, IBM MQ can use it.

See the documentation for your operating system for details of the coded character sets that it supports.

You need to consider application data conversion, format names, and user exits when writing applications that span multiple platforms. See “Writing data-conversion exits” on page 364 for information about invoking and writing data-conversion exits.

Message priorities

You can either set the priority of message to a numeric value, or let the message take the default priority of the queue.

You set the priority of a message (in the *Priority* field of the MQMD structure) when you put the message on a queue. You can set a numeric value for the priority, or you can let the message take the default priority of the queue.

The **MsgDeliverySequence** attribute of the queue determines whether messages on the queue are stored in FIFO (first in, first out) sequence, or in FIFO within priority sequence. If this attribute is set to MQMDS_PRIORITY, messages are enqueued with the priority specified in the *Priority* field of their message descriptors; but if it is set to MQMDS_FIFO, messages are enqueued with the default priority of the queue. Messages of equal priority are stored on the queue in order of arrival.

The **DefPriority** attribute of a queue sets the default priority value for messages being put on that queue. This value is set when the queue is created, but it can be changed afterward. Alias queues, and local definitions of remote queues, can have different default priorities from the base queues to which they resolve. If there is more than one queue definition in the resolution path (see “Name resolution” on page 102), the default priority is taken from the value (at the time of the put operation) of the **DefPriority** attribute of the queue specified in the open command.

The value of the **MaxPriority** attribute of the queue manager is the maximum priority that you can assign to a message processed by that queue manager. You cannot change the value of this attribute. In IBM MQ, the attribute has the value 9; you can create messages having priorities between 0 (the lowest) and 9 (the highest).

Message properties

Use message properties to allow an application to select messages to process, or to retrieve information about a message without accessing MQMD or MQRFH2 headers. They also facilitate communication between IBM MQ and JMS applications.

A message property is data associated with a message, consisting of a textual name and a value of a particular type. Message properties are used by message selectors to filter publications to topics or to selectively get messages from queues. Message properties can be used to include business data or state information without having to store it in the application data. Applications do not have to access data in the MQ Message Descriptor (MQMD) or MQRFH2 headers because fields in these data structures can be accessed as message properties using Message Queue Interface (MQI) function calls.

The use of message properties in IBM MQ mimics the use of properties in JMS. This means that you can set properties in a JMS application and retrieve them in a procedural IBM MQ application, or the other way round. To make a property available to a JMS application, assign it the prefix "usr"; it is then available (without the prefix) as a JMS message user property. For example, the IBM MQ property *usr.myproperty* (a character string) is accessible to a JMS application using the JMS call `message.getStringProperty('myproperty')`. Note that JMS applications are unable to access properties with the prefix "usr" if they contain two or more U+002E (".") characters. A property with no prefix and no U+002E (".") character is treated as if it had the prefix "usr". Conversely, a user property set in a JMS application can be accessed in an IBM MQ application by adding the "usr." prefix to the property name inquired on in an MQINQMP call.

Message properties and message length:

Use the queue manager attribute *MaxPropertiesLength* to control the size of the properties that can flow with any message in an IBM MQ queue manager.

In general, when you use MQSETMP to set properties, the size of a property is the length of the property name in bytes, plus the length of the property value in bytes as passed into the MQSETMP call. It is possible for the character set of the property name and the property value to change during transmission of the message to its destination because these can be converted into Unicode; in this case the size of the property might change.

On an MQPUT or MQPUT1 call, properties of the message do not count toward the length of the message for the queue and the queue manager, but they do count toward the length of the properties as perceived by the queue manager (whether they were set using the message property MQI calls or not).

If the size of the properties exceeds the maximum properties length, the message is rejected with MQRC_PROPERTIES_TOO_BIG. Because the size of the properties is dependent on its representation, you should set the maximum properties length at a gross level.

It is possible for an application to successfully put a message with a buffer that is larger than the value of *MaxMsgLength*, if the buffer includes properties. This is because, even when represented as MQRFH2 elements, message properties do not count toward the length of the message. The MQRFH2 header fields add to the properties length only if one or more folders are contained and every folder in the header contains properties. If one or more folders are contained in the MQRFH2 header and any folder does not contain properties, the MQRFH2 header fields count toward the message length instead.

On an MQGET call, properties of the message do not count toward the length of the message as far as the queue and the queue manager are concerned. However, because the properties are counted separately it is possible that the buffer returned by an MQGET call is larger than the value of the *MaxMsgLength* attribute.

Do not have your applications query the value of *MaxMsgLength* and then allocate a buffer of this size before calling MQGET; instead, allocate a buffer you consider large enough. If the MQGET fails, allocate a buffer guided by the size of the *DataLength* parameter.

The *DataLength* parameter of the MQGET call returns the length in bytes of the application data and any properties returned in the buffer you have provided, if a message handle is not specified in the MQGMO structure.

The *Buffer* parameter of the MQPUT call contains the application message data to be sent and any properties represented in the message data.

When flowing to a queue manager that is earlier than Version 7.0 of the product, properties of the message, except those in the message descriptor, count toward the length of the message. Therefore, you should either raise the value of the *MaxMsgLength* attribute of channels going to a system earlier than Version 7.0 as necessary, to compensate for the fact that more data might be sent for each message. Alternatively, you can lower the queue or queue manager *MaxMsgLength*, so that the overall level of data being sent around the system remains the same.

There is a length limit of 100 MB for message properties, excluding the message descriptor or extension for each message.

The size of a property in its internal representation is the length of the name, plus the size of its value, plus some control data for the property. There is also some control data for the set of properties after one property is added to the message.

Property names:

A property name is a character string. Certain restrictions apply to its length and the set of characters that can be used.

A property name is a case-sensitive character string, limited to +4095 characters unless otherwise restricted by the context. This limit is contained in the MQ_MAX_PROPERTY_NAME_LENGTH constant.

If you exceed this maximum length when using a message property MQI call, the call fails with reason code MQRC_PROPERTY_NAME_LENGTH_ERR.

Because there is no maximum property name length in JMS, it is possible for a JMS application to set a valid JMS property name that is not a valid IBM MQ property name when stored in an MQRFH2 structure.

In this case, when parsed, only the first 4095 characters of the property name are used; the following characters are truncated. This could cause an application using selectors to fail to match a selection string, or to match a string when not expecting to, since more than one property might truncate to the same name. When a property name is truncated, WebSphereMQ issues an error log message.

All property names must follow the rules defined by the Java Language Specification for Java Identifiers, with the exception that Unicode character U+002E (.) is permitted as part of the name - but not the start. The rules for Java Identifiers equate to those contained in the JMS specification for property names.

White space characters and comparison operators are prohibited. Embedded nulls are allowed in a property name but not recommended. If you use embedded nulls, this prevents the use of the MQVS_NULL_TERMINATED constant when used with the MQCHARV structure to specify variable length strings.

Keep property names simple because applications can select messages based on the property names and the conversion between the character set of the name and of the selector might cause the selection to fail unexpectedly.

IBM MQ property names use character U+002E (.) for logical grouping of properties. This divides up the namespace for properties. Properties with the following prefixes, in any mixture of lowercase or uppercase are reserved for use by the product:

- mcd
- jms
- usr
- mq
- sib
- wmq
- Root
- Body
- Properties

A good way to avoid name clashes is to ensure that all applications prefix their message properties with their Internet domain name. For example, if you are developing an application using domain name `ourcompany.com` you could name all properties with the prefix `com.ourcompany`. This naming convention also allows for easy selection of properties; for example, an application can inquire on all message properties starting `com.ourcompany.%`.

See Property name restrictions for further information about the use of property names.

Property name restrictions:

When you name a property, you must observe certain rules.

The following restrictions apply to property names:

1. A property must not begin with the following strings:
 - "JMS" - reserved for use by IBM MQ classes for JMS.
 - "usr.JMS" - not valid.

The only exceptions are the following properties providing synonyms for JMS properties:

Property	Synonym for
JMSCorrelationID	Root .MQMD.CorrelId or jms.Cid
JMSDeliveryMode	Root .MQMD.Persistence or jms.Dlv
JMSDestination	jms.Dst
JMSExpiration	Root .MQMD.Expiry or jms.Exp
JMSMessageID	Root .MQMD.MsgId
JMSPriority	Root .MQMD.Priority or jms.Pri
JMSRedelivered	Root .MQMD.BackoutCount
JMSReplyTo (a string encoded as a URI)	Root .MQMD.ReplyToQ or Root .MQMD.ReplyToQMGr or jms.Rto
JMSTimestamp	Root .MQMD.PutDate or Root .MQMD.PutTime or jms.Tms
JMSType	mcd.Type or mcd.Set or mcd.Fmt
JMSXAppID	Root .MQMD.PutApplName

Property	Synonym for
JMSXDeliveryCount	Root .MQMD.BackoutCount
JMSXGroupID	Root .MQMD.GroupId or jms.Gid
JMSXGroupSeq	Root .MQMD.MsgSeqNumber or jms.Seq
JMSXUserID	Root .MQMD.UserIdentifier

These synonyms allow an MQI application to access JMS properties in a similar fashion to IBM MQ classes for JMS client application. Of these properties, only JMSCorrelationID, JMSReplyTo, JMSType, JMSXGroupID, and JMSXGroupSeq can be set using the MQI.

Note that the JMS_IBM_* properties available from within IBM MQ classes for JMS are not available using the MQI. The fields that the JMS_IBM_* properties reference can be accessed in other ways by MQI applications.

2. A property must not be called, in any mixture of lower or uppercase, "NULL", "TRUE", "FALSE", "NOT", "AND", "OR", "BETWEEN", "LIKE", "IN", "IS" and "ESCAPE". These are the names of SQL keywords used in selection strings.
3. A property name beginning "mq" in any mixture of lowercase or uppercase and not beginning "mq_usr" can contain only one "." character (U+002E). Multiple "." characters are not allowed in properties with those prefixes.
4. Two "." characters must contain other characters in between; you cannot have an empty point in the hierarchy. Similarly a property name cannot end in a "." character.
5. If an application sets the property "a.b" and then the property "a.b.c", it is unclear whether in the hierarchy "b" contains a value or another logical grouping. Such a hierarchy is "mixed content" and this is not supported. Setting a property that causes mixed content is not allowed.

These restrictions are enforced by the validation mechanism as follows:

- Property names are validated when setting a property using the MQSETMP - Set message property call, if validation was requested when the message handle was created. If an attempt to validate a property is undertaken and fails due to an error in the specification of the property name, the completion code is MQCC_FAILED with reason:
 - MQRC_PROPERTY_NAME_ERROR for reasons 1-4.
 - MQRC_MIXED_CONTENT_NOT_ALLOWED for reason 5.
- The names of properties specified directly as MQRFH2 elements are not guaranteed to be validated by the MQPUT call.

Message descriptor fields as properties:

Most message descriptor fields can be treated as properties. The property name is constructed by adding a prefix to the message descriptor field's name.

If an MQI application wants to identify a message property contained in a message descriptor field, for example, in a selector string or using the message property APIs, use the following syntax:

Property name	Message descriptor field
Root.MQMD.Field	Field

Specify *Field* with the same case as for the MQMD structure fields in the C language declaration. For example, the property name Root.MQMD.AccountingToken accesses the AccountingToken field of the message descriptor.

The StrucId and Version fields of the message descriptor are not accessible using the syntax shown.

Message descriptor fields are never represented in an MQRFH2 header as for other properties.

If the message data starts with an MQMDE that is honored by the queue manager, the MQMDE fields can be accessed using the `Root.MQMD.Field` notation described. In this case the MQMDE fields are treated as logically part of the MQMD from a properties perspective. See Overview of MQMDE.

Property data types and values:

A property can be a boolean, a byte string, a character string, or a floating-point or integer number. The property can store any valid value in the range of the data type unless otherwise restricted by the context.

The data type of a property value must be one of the following values:

- MQBOOL
- MQBYTE[]
- MQCHAR[]
- MQFLOAT32
- MQFLOAT64
- MQINT8
- MQINT16
- MQINT32
- MQINT64

A property can exist but have no defined value; it is a null property. A null property is different from a byte property (MQBYTE[]) or character string property (MQCHAR[]) in that it has a defined but empty value, that is, one with a zero-length value.

Byte string is not a valid property data type in JMS or XMS. You are advised not to use byte string properties in the *usr* folder.

Selecting messages from queues

You can select messages from queues using the `MsgId` and `CorrelId` fields on an MQGET call, or by using a `SelectionString` on an MQOPEN or MQSUB call.

Selectors:

A message selector is a variable-length string used by an application to register its interest in only those messages that have properties that satisfy the Structured Query Language (SQL) query that the selection string represents.

Selection using the MQSUB and MQOPEN function calls

You use the *SelectionString*, which is a structure of type MQCHARV, to make selections using the MQSUB and MQOPEN calls.

The *SelectionString* structure is used to pass a variable-length selection string to the queue manager.

The CCSID associated with the selector string is set via the VSCCSID field of the MQCHARV structure. The value used must be a CCSID that is supported for selector strings. See Code page conversion for a list of supported code pages.

Specifying a CCSID for which there is no IBM MQ supported Unicode conversion, results in an error of MQRC_SOURCE_CCSID_ERROR. This error is returned at the time that the selector is presented to the queue manager, that is, on the MQSUB, MQOPEN, or MQPUT1 call.

The default value for the *VSCSID* field is *MQCCSI_APPL*, which indicates that the CCSID of the selection string is equal to the queue manager CCSID, or the client CCSID if connected through a client. The *MQCCSI_APPL* constant can however be overridden by an application redefining it before compiling.

If the *MQCHARV* selector represents a NULL string, no selection takes place for that message consumer and messages are delivered as if a selector had not been used.

The maximum length of a selection string is limited only by what can be described by the *MQCHARV* field *VSLength*.

The SelectionString is returned on the output from an *MQSUB* call using the *MQSO_RESUME* subscribe option, if you have provided a buffer and there is a positive buffer length in *VSBufSize*. If you do not provide a buffer, only the length of the selection string is returned in the *VSLength* field of the *MQCHARV*. If the buffer provided is smaller than the space required to return the field, only *VSBufSize* bytes are returned in the provided buffer.

An application cannot alter a selection string without first closing either the handle to the queue (for *MQOPEN*), or subscription (for *MQSUB*). A new selection string can then be specified on a subsequent *MQOPEN* or *MQSUB* call.

MQOPEN

Use *MQCLOSE* to close the opened handle, then specify a new selection string on a subsequent *MQOPEN* call.

MQSUB

Use *MQCLOSE* to close the returned subscription handle (*hSub*), then specify a new selection string on a subsequent *MQSUB* call.

Figure 3 on page 24 shows the process of selection using the *MQSUB* call.

MQOPEN

(APP 1)

ObjectName = "MyDestQ"
hObj



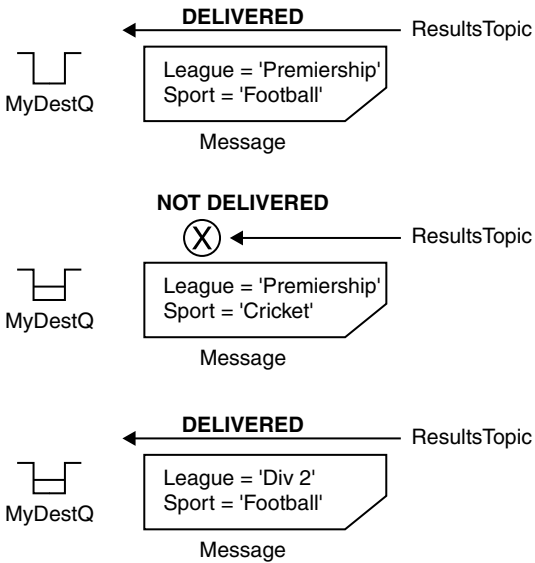
MQSUB

(APP 1)

SelectionString = "Sport = 'Football'"
hObj
TopicString = "ResultsTopic"



ResultsTopic



MQGET

(APP 1) hObj

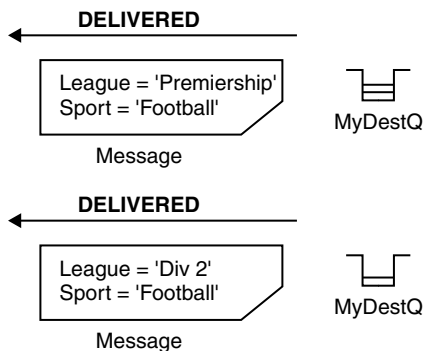


Figure 3. Selection using MQSUB call

A selector can be passed in on the call to MQSUB by using the *SelectionString* field in the MQSD structure. The effect of passing in a selector on the MQSUB is that only those messages published to the topic being subscribed to, that match a supplied selection string, are made available on the destination queue.

Figure 4 on page 25 shows the process of selection using the MQOPEN call.

MQOPEN

(APP 1)

SelectorString = "League = 'Premiership'"
ObjectName = "SportQ"
hObj





SportQ



SportQ




SportQ



MQGET

(APP 1) hObj

NOT DELIVERED




Message



SportQ

DELIVERED



Message



SportQ

MQRC_NO_MSG_AVAILABLE



SportQ

Figure 4. Selection using MQOPEN call

A selector can be passed in on the call to MQOPEN by using the *SelectionString* field in the MQOD structure. The effect of passing in a selector on the MQOPEN call is that only those messages on the opened queue, that match a selector, are delivered to the message consumer.

The main use for the selector on the MQOPEN call is for the point-to-point case where an application can elect to receive only those messages on a queue that match a selector. The previous example shows a simple scenario where two messages are put to a queue opened by MQOPEN but only one is received by the application getting it, as it is the only one that matches a selector.

Note that subsequent MQGET calls result in MQRC_NO_MSG_AVAILABLE as no further messages exist on the queue that match the given selector.

Related concepts:

“Selection string rules and restrictions” on page 32

Familiarize yourself with these rules about how selection strings are interpreted and character restrictions to avoid potential problems when using selectors.

Selection behavior:

Overview of IBM MQ selection behavior.

The fields in an MQMDE structure are considered to be the message properties for the corresponding message descriptor properties if the MQMD:

- Has format MQFMT_MD_EXTENSION
- Is immediately followed by a valid MQMDE structure
- Is version one or contains the default version two fields only

It is possible for a selection string to resolve to either TRUE or FALSE before any matching against message properties takes place. For example, it might be the case if the selection string is set to "TRUE <> FALSE". Such early evaluation is guaranteed to take place only when there are no message property references in the selection string.

If a selection string resolves to TRUE before any message properties are considered, all messages published to the topic subscribed to by the consumer are delivered. If a selection string resolves to FALSE before any message properties are considered, a reason code of MQRC_SELECTOR_ALWAYS_FALSE, and completion code MQCC_FAILED are returned on the function call that presented the selector.

Even if a message contains no message properties (other than header properties) then it can still be eligible for selection. If a selection string references a message property that does not exist, this property is assumed to have the value of NULL or 'Unknown'.

For example, a message might still satisfy a selection string like 'Color IS NULL', where 'Color' does not exist as a message property in the message.

Selection can be performed only on the properties that are associated with a message, not the message itself, unless an extended message selection provider is available. Selection can be performed on the message payload only if an extended message selection provider is available.

Each message property has a type associated with it. When you perform a selection, you must ensure that the values used in expressions to test message properties are of the correct type. If a type mismatch occurs, the expression in question resolves to FALSE.

It is your responsibility to ensure that the selection string and message properties use compatible types.

Selection criteria continue to be applied on behalf of inactive durable subscribers, so that only messages that match the selection string that was originally supplied are kept.

Selection strings are non-alterable when a durable subscription is resumed with alter (MQSO_ALTER). If a different selection string is presented when a durable subscriber resumes activity, then MQRC_SELECTOR_NOT_ALTERABLE is returned to the application.

Applications receive a return code of MQRC_NO_MSG_AVAILABLE if there is no message on a queue that meets the selection criteria.

If an application has specified a selection string containing property values then only those messages that contain matching properties are eligible for selection. For example, a subscriber specifies a selection string

of "a = 3" and a message is published containing no properties, or properties where 'a' does not exist or is not equal to 3. The subscriber does not receive that message to its destination queue.

Messaging performance

Selecting messages from a queue requires IBM MQ to sequentially inspect each message on the queue. Messages are inspected until a message is found that matches the selection criteria or there are no more messages to examine. Therefore, messaging performance suffers if message selection is used on deep queues.

To optimize message selection on deep queues when selection is based on JMSCorrelationID or JMSMessageID, use a selection string of the form JMSCorrelationID = ... or JMSMessageID = ... and reference only one property.

This method offers a significant improvement in performance for selection on JMSCorrelationID and offers a marginal performance improvement for JMSMessageID.

Using complex selectors

Selectors can contain many components, for example:

a and b or c and d or e and f or g and h or i and j... or y and z

Use of such complex selectors can have serious performance implications and excessive resource requirements. As such, IBM MQ will protect the system by failing to process overly complex selectors that could result in a system resource shortage. Protection can occur on selection strings that contain more than 100 tests, or when IBM MQ detects that the limit on the size of the operating system stack is being approached. You should thoroughly try and test the use of selection strings with many components, on the appropriate platforms, to ensure that the protection limits are not reached.

The performance and complexity of selectors can be improved by simplifying them using additional parenthesis to combine components. For example:

(a and b or c and d) or (e and f or g and h) or (i and j) ...

Related concepts:

"Selection string rules and restrictions" on page 32

Familiarize yourself with these rules about how selection strings are interpreted and character restrictions to avoid potential problems when using selectors.

Message selector syntax:

An IBM MQ message selector is a string with syntax that is based on a subset of the SQL92 conditional expression syntax.

The order in which a message selector is evaluated is from left to right within a precedence level. You can use parentheses to change this order. Predefined selector literals and operator names are written here in uppercase; however, they are not case-sensitive.

IBM MQ verifies the syntactic correctness of a message selector at the time it is presented. If the syntax of the selection string is incorrect or a property name is not valid, and an extended message selection provider is not available, MQRC_SELECTION_NOT_AVAILABLE is returned to the application. If the syntax of the selection string is incorrect or a property name is not valid when a subscription is resumed, a MQRC_SELECTOR_SYNTAX_ERROR is returned to the application. If property name validation was disabled when the property was set (by setting MQCMHO_NONE instead of MQCMHO_VALIDATE) and an application subsequently puts a message with an invalid property name, this message is never selected.

A selector can contain:

- Literals:
 - String literals are enclosed in single quotation marks. Two consecutive single quotation marks represent a single quotation mark. Examples are 'literal' and 'literal's'. Like Java string literals, these use the Unicode character encoding. You cannot use double quotation marks to enclose a string literal. Any sequence of bytes can be used between the single quotation marks.
 - A byte string is one or more pair of hexadecimal characters enclosed in double quotation marks and prefixed by 0x. Examples are "0x2F1C" or "0XD43A". The length of a byte string must be at least one byte. If a selector byte string is matched to a message property of type MQTYPE_BYTE_STRING, no special action is taken on leading or trailing zero. The bytes are treated as another character. Endianness is also not considered. The length of both selector and property byte strings must be equal, and the sequence of bytes must be the same.

Examples of byte string selections (assume *myBytes* = 0AFC23) which match are:

- "myBytes = "0x0AFC23"" = TRUE

The following string selections do not match:

- "myBytes = "0xAFC23"" = MQRC_SELECTOR_SYNTAX_ERROR (because number of bytes is not multiple of two)
- "myBytes = "0x0AFC2300"" = FALSE (because the trailing zero is significant in the comparison)
- "myBytes = "0x000AFC23"" = FALSE (because leading zero is significant in the comparison)
- "myBytes = "0x23FC0A"" = FALSE (because endianness is not considered)

- Hex numbers begin with a zero, followed by an uppercase or lowercase x. The remainder of the literal contains one or more valid hex characters. Examples are 0xA, 0xAF, 0X2020.
- A leading zero followed by one or more digits in the range 0-7 is always interpreted as being the start of an octal number. You cannot represent a zero-prefixed decimal number like this, for example, 09 returns a syntax error because 9 is not a valid octal digit. Examples of octal numbers are 0177, 0713.
- An exact numeric literal is a numeric value without a decimal point, such as 57, -957, and +62. An exact numeric literal can have a trailing uppercase or lowercase L; this does not affect how the number is stored or interpreted. IBM MQ supports exact numerals in the range -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807.
- An approximate numeric literal is a numeric value in scientific notation, such as 7E3 or -57.9E2, or a numeric value with a decimal, such as 7., -95.7, or +6.2. IBM MQ supports numbers in the range -1.797693134862315E+308 to 1.797693134862315E+308.

The significand should follow an optional sign character (+ or -). The significand should be either an integer or a fraction. A fractional part of the significand need not have a leading digit.

An uppercase or lowercase E indicates the start of an optional exponent. The exponent has a decimal radix and the number part of the exponent can be prefixed by an optional sign character.

Approximate numeric literals can be terminated by an F or a D character (not case-sensitive). This syntax exists to support the cross-language method of tagging single or double precision numbers. These characters are optional and do not affect how an approximate numeric literal is stored or processed. These numbers are always stored and processed using double-precision.

- The boolean literals TRUE and FALSE.

Note: Non-finite IEEE-754 representations such as NaN, +Infinity, -Infinity are not supported in selection strings. It is therefore not possible to use these values as operands in an expression. Negative zero is treated the same as positive zero for mathematical operations.

- Identifiers:

An identifier is a variable-length character sequence that must begin with a valid identifier start character, followed by zero or more valid identifier part characters. The rules for identifier names are the same as those for message property names, see "Property names" on page 19 and "Property name restrictions" on page 20 for more information.

Note: Selection can be performed on the message payload only if an extended message selection provider is available.

Identifiers are either header field references or property references. The type of a property value in a message selector must correspond to the type used to set the property, although numeric promotion is performed where possible. If a type mismatch occurs then the result of the expression is FALSE. If a property that does not exist in a message is referenced, its value is NULL.

Type conversions that apply to the get methods for properties do not apply when a property is used in a message selector expression. For example, if you set a property as a string value and then use a selector to query it as a numeric value, the expression returns FALSE.

JMS field and property names that map to property names or MQMD field names are also valid identifiers in a selection string. IBM MQ maps the recognized JMS field and property names to the message property values. See “Message selectors in JMS” on page 875 for more information. As an example, the selection string "JMSPriority >=" selects on the Pri property found in the jms folder of the current message.

- Overflow/underflow:

For both decimal and approximate numeric numbers, the following conditions are undefined:

- Specifying a number that is out of the defined range
- Specifying an arithmetic expression which would cause overflow or underflow

No checks are performed for these conditions.

- White space:

Defined as a space, form-feed, newline, carriage return, horizontal tab, or vertical tab. The following Unicode characters are recognized as white space:

- \u0009 to \u000D
- \u0020
- \u001C
- \u001D
- \u001E
- \u001F
- \u1680
- \u180E
- \u2000 to \u200A
- \u2028
- \u2029
- \u202F
- \u205F
- \u3000

- Expressions:

- A selector is a conditional expression. A selector that evaluates to true matches; a selector that evaluates to false or unknown does not match.
- Arithmetic expressions are composed of themselves, arithmetic operations, identifiers (identifier value is treated as a numeric literal), and numeric literals.
- Conditional expressions are composed of themselves, comparison operations, and logical operations.

- Standard bracketing (), to set the order in which expressions are evaluated, is supported.

- Logical operators in precedence order: NOT, AND, OR.

- Comparison operators: =, >, >=, <, <=, <> (not equal).

- Two byte strings are equal only if the strings are of the same length and the sequence of bytes is equal.

- Only values of the same type can be compared. One exception is that it is valid to compare exact numeric values and approximate numeric values, (the type conversion required is defined by the rules of Java numeric promotion). If there is an attempt to compare different types, the selector is always false.
- String and boolean comparison is restricted to = and <>. Two strings are equal only if they contain the same sequence of characters.
- Arithmetic operators in precedence order:
 - +, - unary.
 - * multiplication, and / division.
 - + addition, and - subtraction.
 - Arithmetic operations on a NULL value are not supported. If they are attempted, the complete selector is always false.
 - Arithmetic operations must use Java numeric promotion.
- arithmetic-expr1 [NOT] BETWEEN arithmetic-expr2 and arithmetic-expr3 comparison operator:
 - Age BETWEEN 15 and 19 is equivalent to age >= 15 AND age <= 19.
 - Age NOT BETWEEN 15 and 19 is equivalent to age < 15 OR age > 19.
 - If any of the expressions of a BETWEEN operation are NULL, the value of the operation is false. If any of the expressions of a NOT BETWEEN operation are NULL, the value of the operation is true.
- identifier [NOT] IN (string-literal1, string-literal2,...) comparison operator where identifier has a String or NULL value.
 - Country IN ('UK', 'US', 'France') is true for 'UK' and false for 'Peru'. It is equivalent to the expression (Country = 'UK') OR (Country = 'US') OR (Country = 'France').
 - Country NOT IN ('UK', 'US', 'France') is false for 'UK' and true for 'Peru'. It is equivalent to the expression NOT ((Country = 'UK') OR (Country = 'US') OR (Country = 'France')).
 - If the identifier of an IN or NOT IN operation is NULL, the value of the operation is unknown.
- identifier [NOT] LIKE *pattern-value* [ESCAPE *escape-character*] comparison operator, where identifier has a string value. *pattern-value* is a string literal, where _ stands for any single character and % stands for any sequence of characters (including the empty sequence). All other characters stand for themselves. The optional *escape-character* is a single character string literal that is used to escape the special meaning of the _ and % in *pattern-value*. The LIKE operator must be used only to compare two string values.
 - phone LIKE '12%3' is true for 123 and 12993 and false for 1234.
 - word LIKE 'l_se' is true for lose and false for loose.
 - underscored LIKE '_%' ESCAPE '\\' is true for _foo and false for bar.
 - phone NOT LIKE '12%3' is false for 123 and 12993 and true for 1234.
 - If the identifier of a LIKE or NOT LIKE operation is NULL, the value of the operation is unknown.

Note: The LIKE operator must be used to compare two string values. The value of Root.MQMD.CorrelId is a 24-byte byte array, not a character string. The selector string Root.MQMD.CorrelId LIKE 'ABC%' is accepted by the parser as syntactically valid, but it is evaluated to false. When you are comparing a byte array with a character string, LIKE therefore cannot be used.

- identifier IS NULL comparison operator tests for a NULL header field value, or a missing property value.
- identifier IS NOT NULL comparison operator tests for the existence of a non-null header field value or a property value.
- Null values

The evaluation of selector expressions that contain NULL values is defined by SQL 92 NULL semantics, in summary:

- SQL treats a NULL value as unknown.

- Comparison or arithmetic with an unknown value always yields an unknown value.
- The IS NULL and IS NOT NULL operators convert an unknown value into TRUE and FALSE values.

The boolean operators use three-valued logic (T=TRUE, F=FALSE, U=UNKNOWN)

Table 1. Boolean operator outcome when logic is A AND B

Operator A	Operator B	Outcome (A AND B)
T	F	F
T	U	U
T	T	T
F	T	F
F	U	F
F	F	F
U	T	U
U	U	U
U	F	F

Table 2. Boolean operator outcome when logic is A OR B

Operator A	Operator B	Outcome (A OR B)
T	F	T
T	U	T
T	T	T
F	T	T
F	U	U
F	F	F
U	T	T
U	U	U
U	F	U

Table 3. Boolean operator outcome when logic is NOT A

Operator A	Outcome (NOT A)
T	F
F	T
U	U

The following message selector selects messages with a message type of car, color of blue, and weight greater than 2500 lbs:

```
"JMSType = 'car' AND color = 'blue' AND weight > 2500"
```

Although SQL supports fixed decimal comparison and arithmetic, message selectors do not. This is why exact numeric literals are restricted to those without a decimal. It is also why there are numerics with a decimal as an alternative representation for an approximate numeric value.

SQL comments are not supported.

Related concepts:

“Message properties” on page 18

Use message properties to allow an application to select messages to process, or to retrieve information about a message without accessing MQMD or MQRFH2 headers. They also facilitate communication between IBM MQ and JMS applications.

Related information:

MsgHandle

MQBUFMH - Convert buffer into message handle

Selection string rules and restrictions:

Familiarize yourself with these rules about how selection strings are interpreted and character restrictions to avoid potential problems when using selectors.

- Message selection for publish/subscribe messaging occurs on the message as sent by the publisher. See Selection strings.
- Equivalence is tested using a single equals character; for example, a = b is correct, whereas a == b is incorrect.
- An operator used by many programming languages to represent 'not equal to' is !=. This representation is not a valid synonym for <> ; for example, a <> b is valid, whereas a != b is not valid.
- Single quotation marks are recognized only if the ' (U+0027) character is used. Similarly, double quotation marks, valid only when used to enclose byte strings, must use the " (U+0022) character.
- The symbols &, &&, | and || are not synonyms for logical conjunction/disjunction; for example, a && b must be specified as a AND b.
- The wildcard characters * and ? are not synonyms for % and _.
- Selectors containing compound expressions such as 20 < b < 30 are not valid. The parser evaluates operators that have the same precedence from left to right. The example would therefore become (20 < b) < 30, which does not make sense. Instead the expression must be written as (b > 20) AND (b < 30).
- Byte strings must be enclosed in double quotation marks; if single quotation marks are used, the byte string is taken to be a string literal. The number of characters (not the number that the characters represent) following the 0x must be a multiple of two.
- The keyword IS is not a synonym for the equals character. Thus the selection strings a IS 3 and b IS 'red' are not valid. The IS keyword exists only to support IS NULL and IS NOT NULL cases.

Related concepts:

“Selection behavior” on page 26


Overview of IBM MQ selection behavior.

Related information:

Selection strings

UTF-8 and Unicode considerations when using message selectors:

Characters, not enclosed in single quotation marks, that make up the reserved keywords of a selection string must be entered in Basic Latin Unicode (ranging from character U+0000 to U+0007F). It is not valid to use other code point representations of alphanumeric characters. For example, the number 1 must be expressed as U+0031 in Unicode, it is not valid to use the Fullwidth Digit equivalent U+FF11 or the Arabic equivalent U+0661.

Message property names can be specified using any valid sequence of Unicode characters. Message property names contained within selection strings that are encoded in UTF-8 will be validated even if they contain multi-byte characters. Validation of multi-byte UTF-8 is strict and you must ensure that valid UTF-8 sequences are used for message property names.  Characters beyond the Unicode

Basic Multilingual Plane (those above U+FFFF), represented in UTF-16 by surrogate code points (X'D800' through X'DFFF'), or four bytes in UTF-8, are not supported in message property names.

No extra processing is performed on property names or values when comparing for equality. This means for example that no pre/de-composition takes place and ligatures are not given any special meaning. For example, the pre-composed umlaut character U+00FC is not considered to be equivalent to U+0075 + U+0308 and the character sequence ff is not considered to be equivalent to the Unicode U+FB00 (LATIN SMALL LIGATURE FF)

Property data enclosed in single quotation marks can be represented by any sequence of bytes and is not validated.

Selecting on the content of a message:

It is possible to subscribe based on a selection of message payload content (also known as content filtering), but the decision about which messages should be delivered to such a subscription cannot be performed directly by IBM MQ; instead an extended message selection provider, for example IBM Integration Bus, is required to process the messages.

When an application publishes on a topic string, where one or more subscribers have a selection string selecting on the content of the message, IBM MQ will request that the extended message selection provider parse the publication and inform IBM MQ whether the publication matches the selection criteria specified by each subscriber with a content filter.

If the extended message selection provider determines that the publication matches the subscriber's selection string, the message will continue to be delivered to the subscriber.

If the extended message selection provider determines that the publication does not match, the message is not delivered to the subscriber. This might cause the MQPUT or MQPUT1 call to fail with reason code MQRC_PUBLICATION_FAILURE. If the extended message selection provider is unable to parse the publication, reason code MQRC_CONTENT_ERROR is returned and the MQPUT or MQPUT1 call fails.

If the extended message selection provider is unavailable or is unable to determine whether the subscriber should receive the publication, reason code MQRC_SELECTION_NOT_AVAILABLE is returned and the MQPUT or MQPUT1 call fails.

When a subscription is being created with a content filter and the extended message selection provider is not available, the MQSUB call fails with reason code MQRC_SELECTION_NOT_AVAILABLE. If a subscription with a content filter is being resumed and the extended message selection provider is not available, the MQSUB call returns a warning of MQRC_SELECTION_NOT_AVAILABLE, but the subscription is allowed to be resumed.

Related information:

Selection strings

Asynchronous consumption of IBM MQ messages

Asynchronous consumption uses a set of Message Queue Interface (MQI) extensions, the MQI calls MQCB and MQCTL, which allow an MQI application to be written to consume messages from a set of queues. Messages are delivered to the application by invoking a 'unit of code', identified by the application passing either the message, or a token representing the message.

In the most straightforward of application environments, the unit of code is defined by a function pointer, however in other environments the unit of code can be defined by a program or module name.

In asynchronous consumption of messages, the following terms are used:

Message consumer

A programming construct that allows you to define a program, or function, to be invoked with a message when one which matches the applications requirement becomes available.

Event handler

A programming construct that allows you to define a program or function to invoke when an asynchronous event, such as queue manager quiescing, occurs.

Callback

A generic term used to refer to either a Message Consumer or an Event Handler routine.

Asynchronous consumption can simplify the design and implementation of new applications, especially those that process multiple input queues or subscriptions. However, if you are using more than one input queue and you are processing messages in priority sequence, priority sequence is observed independently within each queue: you might get low-priority messages from one queue ahead of high-priority messages from another. Message order across multiple queues is not guaranteed. Also note that if you use API exits, you might need to change them to include the MQCB and MQCTL calls.

The following illustrations give an example of how you can use this function.

Figure 5 on page 35 shows a multithreaded application consuming messages from two queues. The example shows all of the messages being delivered to a single function.

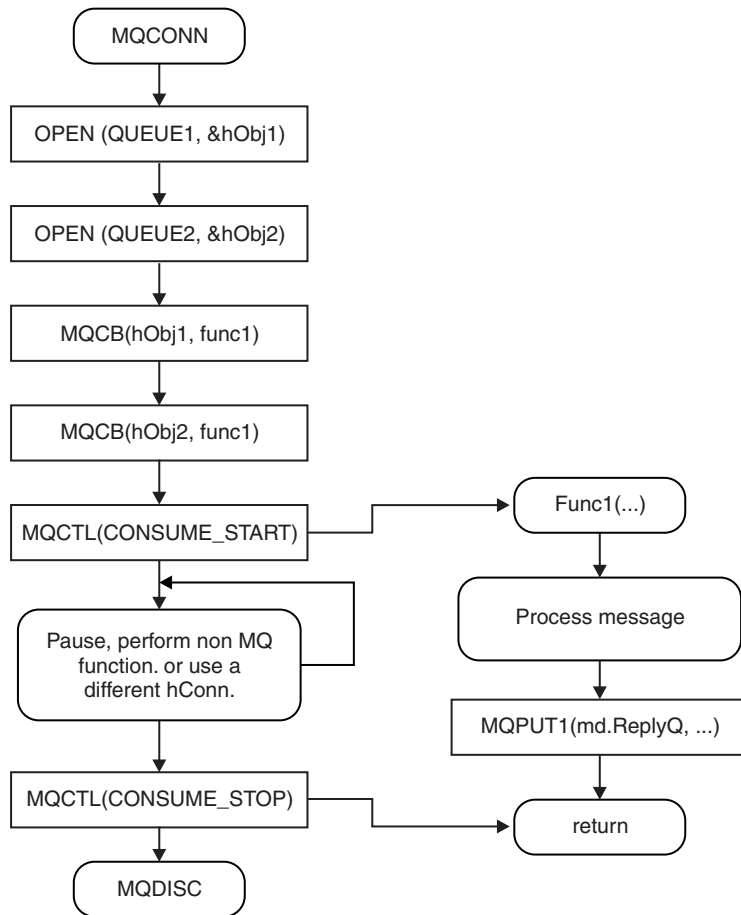


Figure 5. Standard Message Driven application consuming from two queues

z/OS On z/OS, the main control thread must issue an MQDISC call before ending. This allows any callback threads to end and release system resources.

Figure 6 on page 36 This sample flow shows a single threaded application consuming messages from two queues. The example shows all of the messages being delivered to a single function.

The difference from the asynchronous case is that control does not return to the issuer of MQCTL until all of the consumers have deactivated themselves; that is one consumer has issued an MQCTL STOP request or the queue manager quiesces.

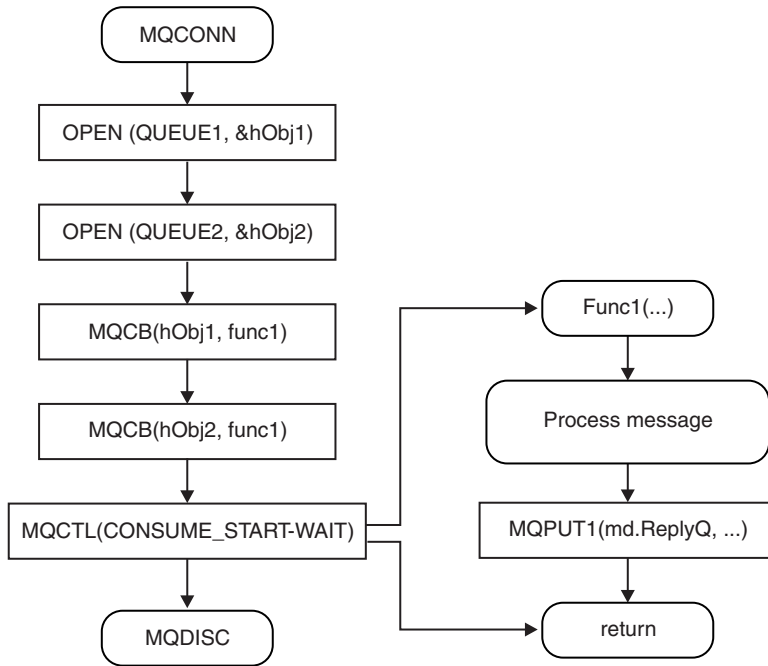


Figure 6. Single Threaded Message Driven application consuming from two queues

Message groups

Messages can occur within groups to allow ordering of messages.

Message groups allow multiple messages to be marked as related to one another, and a logical order to be applied to the group (see “Logical and physical ordering” on page 134). On Multiplatforms, message segmentation enables large messages to be broken up into smaller segments. You cannot use grouped or segmented messages when putting to a topic.

The hierarchy within a group is as follows:

Group This is the highest level in the hierarchy and is identified by a *GroupId*. It consists of one or more messages that contain the same *GroupId*. These messages can be stored anywhere on the queue.

Note: The term *message* is used here to denote one item on a queue, such as would be returned by a single MQGET that does not specify MQGMO_COMPLETE_MSG.

Figure 7 shows a group of logical messages:




Figure 7. Group of logical messages

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. For more information on the BIND_ON_GROUP option, see Handling message affinities.

Logical message

Logical messages within a group are identified by the *GroupId* and *MsgSeqNumber* fields. The *MsgSeqNumber* starts at 1 for the first message within a group, and if a message is not in a group, the value of the field is 1.

Use logical messages within a group to:

- Ensure ordering (if this is not guaranteed under the circumstances in which the message is transmitted).
- Allow applications to group similar messages (for example, those that must all be processed by the same server instance).

Each message within a group consists of one physical message, unless it is split into segments. Each message is logically a separate message, and only the *GroupId* and *MsgSeqNumber* fields in the MQMD need to bear any relationship to other messages in the group. Other fields in the MQMD are independent; some might be identical for all messages in the group whereas others might be different. For example, messages in a group can have different format names, CCSIDs, and encodings.

Segment

Segments are used to handle messages that are too large for either the putting or getting application or the queue manager (including intervening queue managers through which the message passes). For more information, see “Message segmentation” on page 153.

An individual message is broken down into smaller messages called *segments*. A segment of a message is identified by the *GroupId*, *MsgSeqNumber*, and *Offset* fields. The *Offset* field starts at zero for the first segment within a message.

Each segment consists of one physical message that might belong to a group (Figure 8 shows an example of messages within a group). A segment is logically part of a single message, so only the *MsgId*, *Offset*, and *MsgFlags* fields in the MQMD should differ between separate segments of the same message. If a segment fails to arrive, reason code MQRC_INCOMPLETE_GROUP or MQRC_INCOMPLETE_MSG is returned as appropriate.

Figure 8 shows a group of logical messages, some of which are segmented:




Figure 8. Segmented messages

► z/OS Segmentation is not supported on IBM MQ for z/OS.

You cannot use segmented or grouped messages with Publish/Subscribe.

Related concepts:

“Message segmentation” on page 153

Use this information to learn about segmenting messages. This feature is not supported on IBM MQ for z/OS or by applications using IBM MQ classes for JMS.

Related reference:

“Logical and physical ordering” on page 134


Messages on queues can occur (within each priority level) in *physical* or *logical* order.

Related information:

MQMD - Message descriptor



Message persistence

Persistent messages are written to logs and queue data files. If a queue manager is restarted after a failure, it recovers these persistent messages as necessary from the logged data. Messages that are not persistent are discarded if a queue manager stops, whether the stoppage is as a result of an operator command or because of the failure of some part of your system.

 Nonpersistent messages stored in a coupling facility (CF) on z/OS are an exception to this. They persist as long as the CF remains available.

When you create a message, if you initialize the message descriptor (MQMD) using the defaults, the persistence for the message is taken from the **DefPersistence** attribute of the queue specified in the MQOPEN command. Alternatively, you can set the persistence of the message using the *Persistence* field of the MQMD structure to define the message as persistent or nonpersistent.

The performance of your application is affected when you use persistent messages; the extent of the effect depends on the performance characteristics of the machine's I/O subsystem and how you use the sync point options on each platform:

- A persistent message, outside the current unit of work, is written to disk on every put and get operation. See “Committing and backing out units of work” on page 222.
-   For all platforms except IBM i, a persistent message within the current unit of work is logged only when the unit of work is committed, and the unit of work can contain many queue operations.

Nonpersistent messages can be used for fast messaging. See Safety of messages for further information about fast messages.

Note: A combination of writing persistent messages within a unit of work, and writing persistent messages outside a unit or work, can cause potentially severe performance problems for your applications. This is particularly true when the same target queue is used for both operations.

Messages that fail to be delivered

When a queue manager cannot put a message on a queue, you have various options.

You can:

- Attempt to put the message on the queue again.
- Request that the message is returned to the sender.
- Put the message on the dead-letter queue.

See “Handling procedural program errors” on page 436 for more information.

Messages that are backed out

When processing messages from a queue under the control of a unit of work, the unit of work can consist of one or more messages. If a backout occurs, the messages that were retrieved from the queue are reinstated on the queue, and they can be processed again in another unit of work. If the processing of a particular message is causing the problem, the unit of work is backed out again. This can cause a processing loop. Messages that were put to a queue are removed from the queue.

An application can detect messages that are caught up in such a loop by testing the *BackoutCount* field of MQMD. The application can either correct the situation, or issue a warning to an operator.

Multi The backout count always survives restarts of the queue manager. Any change to the **HardenGetBackout** attribute is ignored.

z/OS For shared queues, the backout count always survives restarts of the queue manager. For all other configurations on z/OS, to ensure that the backout count for private queues survives restarts of the queue manager, set the *HardenGetBackout* attribute to MQQA_BACKOUT_HARDENED; otherwise, if the queue manager has to restart, it does not maintain an accurate backout count for each message. Setting the attribute this way adds the cost of extra processing.

For more information on committing and backing out messages, see “Committing and backing out units of work” on page 222.

Reply-to queue and queue manager

There are occasions when you might receive messages in response to a message you send:

- A reply message in response to a request message
- A report message about an unexpected event or expiry
- A report message about a COA (Confirmation Of Arrival) or a COD (Confirmation Of Delivery) event
- A report message about a PAN (Positive Action Notification) or a NAN (Negative Action Notification) event

Using the MQMD structure, specify the name of the queue to which you want reply and report messages sent, in the *ReplyToQ* field. Specify the name of the queue manager that owns the reply-to queue in the *ReplyToQMgr* field.

If you leave the *ReplyToQMgr* field blank, the queue manager sets the contents of the following fields in the message descriptor on the queue:

ReplyToQ

If *ReplyToQ* is a local definition of a remote queue, the *ReplyToQ* field is set to the name of the remote queue; otherwise this field is not changed.

ReplyToQMgr

If *ReplyToQ* is a local definition of a remote queue, the *ReplyToQMgr* field is set to the name of the

queue manager that owns the remote queue; otherwise the *ReplyToQMgr* field is set to the name of the queue manager to which your application is connected.

Note: You can request that a queue manager makes more than one attempt to deliver a message, and you can request that the message is discarded if it fails. If the message, after failing to be delivered, is not to be discarded, the remote queue manager puts the message on its dead-letter (undelivered message) queue (see “Using the dead-letter (undelivered message) queue” on page 440).

Message context

Message context information allows the application that retrieves the message to find out about the originator of the message.

The retrieving application might want to:

- Check that the sending application has the correct level of authority
- Perform some accounting function so that it can charge the sending application for any work that it has to perform
- Keep an audit trail of all the messages that it has worked with

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. For more information about how to specify context information, see “Controlling message context information” on page 118.

The user context is used by the queue manager when generating the following types of report message:

- Confirm on delivery
- Expiry

When these report messages are generated, the user context is checked for +put and +passid authority on the destination of the report. Where the user context has insufficient authority, the report message is placed on the dead-letter queue if one has been defined. Where there is no dead-letter queue, the report message is discarded.

All context information is stored in the context fields of the message descriptor. The type of information falls into identity, origin, and user context information.

Identity context

Identity context information identifies the user of the application that first put the message on a queue. Suitably authorized applications can set the following fields:

- The queue manager fills the *UserIdentifier* field with a name that identifies the user; the way that the queue manager can do this depends on the environment in which the application is running.
- The queue manager fills the *AccountingToken* field with a token or number that it determined from the application that put the message.
- Applications can use the *ApplIdentityData* field for any extra information that they want to include about the user (for example, an encrypted password).

A Windows systems security identifier (SID) is stored in the *AccountingToken* field when a message is created under IBM MQ for Windows. The SID can be used to supplement the *UserIdentifier* field and to establish the credentials of a user.

For information about how the queue manager fills the *UserIdentifier* and *AccountingToken* fields, see the descriptions of these fields in *UserIdentifier* and *AccountingToken*.

Applications that pass messages from one queue manager to another should also pass on the identity context information so that other applications know the identity of the originator of the message.

Origin context

Origin context information describes the application that put the message on the queue on which the message is currently stored. The message descriptor contains the following fields for origin context information:

- *PutApplType* defines the type of application that put the message (for example, a CICS transaction).
- *PutApplName* defines the name of the application that put the message (for example, the name of a job or transaction).
- *PutDate* defines the date on which the message was put on the queue.
- *PutTime* defines the time at which the message was put on the queue.
- *ApplOriginData* defines any extra information that an application wants to include about the origin of the message. For example, it could be set by suitably authorized applications to indicate whether the identity data is trusted.

Origin context information is typically supplied by the queue manager. Greenwich Mean Time (GMT) is used for the *PutDate* and *PutTime* fields. See the descriptions of these fields in *PutDate* and *PutTime*.

An application with enough authority can provide its own context. This allows accounting information to be preserved when a single user has a different user ID on each of the systems that process a message that they have originated.

IBM MQ objects

This information provides details on IBM MQ objects which include: queue managers, queue-sharing groups, queues, administrative topic objects, namelists, process definitions, authentication information objects, channels, storage classes, listeners, and services.

Queue managers define the properties (known as attributes) of these objects. The values of these attributes affect the way in which IBM MQ processes these objects. From your applications, you use the Message Queue Interface (MQI) to control these objects. Objects are identified by an *object descriptor* (MQOD) when addressed from a program.

When you use IBM MQ commands to define, alter, or delete objects, for example, the queue manager checks that you have the required level of authority to perform these operations. Similarly, when an application uses the MQOPEN call to open an object, the queue manager checks that the application has the required level of authority before it allows access to that object. The checks are made on the name of the object being opened.

Related concepts:

“Controlling message context information” on page 118

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. You can use the options field in the MQPMO structure to control context information.

Related reference:

“MQOPEN options relating to message context” on page 108

If you want to be able to associate context information with a message when you put it on a queue, you must use one of the message context options when you open the queue.

Preparing and running Microsoft Transaction Server applications

To prepare an MTS application to run as an IBM MQ MQI client application, follow these instructions as appropriate for your environment.

For general information about how to develop Microsoft Transaction Server (MTS) applications that access IBM MQ resources, see the section on MTS in the IBM MQ Help Center.

To prepare an MTS application to run as an IBM MQ MQI client application, do one of the following for each component of the application:

- If the component uses the C language bindings for the MQI, follow the instructions in “Preparing C programs in Windows” on page 414 but link the component with the library mqicxa.lib instead of mqic.lib.
- If the component uses the IBM MQ C++ classes, follow the instructions in “Building C++ programs on Windows” on page 820 but link the component with the library imqx23vn.lib instead of imqc23vn.lib.
- If the component uses the Visual Basic language bindings for the MQI, follow the instructions in the “Preparing Visual Basic programs in Windows” on page 417 but when you define the Visual Basic project, type MqType=3 in the **Conditional Compilation Arguments** field.
- If the component uses the IBM MQ Automation Classes for ActiveX (MQAX), define an environment variable, GMQ_MQ_LIB, with the value mqic32xa.dll.

You can define the environment variable from within your application, or you can define it so that its scope is system wide. However, defining it as system wide can cause any existing MQAX application, that does not define the environment variable from within the application, to behave incorrectly.

Using IBM MQ with WebSphere Application Server

Use this topic to understand the use of IBM MQ with WebSphere Application Server.

Applications that are written in Java that are running under WebSphere Application Server can use the Java Messaging Service (JMS) specification to perform messaging. Messaging in this environment can be provided by an IBM MQ queue manager.

A benefit of using an IBM MQ queue manager is that connecting JMS applications can participate fully in the functionality of an IBM MQ network, which allows the applications to exchange messages with queue managers that are running on a multitude of platforms.

Applications can use either the *client transport* or *bindings transport* for the queue connection factory object. For *bindings transport*, the queue manager must exist locally to the application that requires a connection.

By default, JMS messages that are held on IBM MQ queues use an MQRFH2 header to hold some of the JMS message header information. Many legacy IBM MQ applications cannot process messages with these headers, and require their own characteristic headers, for example the MQCIH for CICS Bridge, or MQWIH for IBM MQ Workflow applications. For more information about these special considerations, see Mapping JMS messages onto IBM MQ messages.

Design considerations for IBM MQ applications

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

When designing an IBM MQ application consider the following questions and options:

Type of application

What is the purpose of your application? See the following links for information about that different types of application you can develop:

- Server
- Client
- Publish/subscribe
- Web services
- User exits, API exits, and installable services

Additionally, you can also write your own applications to automate administration of IBM MQ. For more information, see *The IBM MQ Administration Interface (MQAI)* and *Automating administration tasks*.

Programming language

IBM MQ supports a number of procedural and object-oriented programming languages for writing applications. For more information see, “Developing applications for IBM MQ” on page 1.

Applications for more than one platform

Will your application run on more than one platform? Do you have a strategy to move to a different platform from the one that you use today? If the answer to either of these questions is yes, ensure that you code your programs for platform independence.

For example if you are using C, code in ANSI standard C. Use a standard C library function rather than an equivalent platform-specific function even if the platform-specific function is faster or more efficient. The exception is when efficiency in the code is paramount, when you should code for both situations using `#ifdef`. For example:

```
#ifdef _AIX
    AIX specific code
#else
    generic code
#endif
```

Types of queues

Do you want to create a queue each time that you need one, or do you want to use queues that have already been set up? Do you want to delete a queue when you have finished using it, or is it going to be used again? Do you want to use alias queues for application independence? To see what types of queues are supported, refer to *Queues*.

Using shared queues, queue-sharing groups, and queue-sharing group clusters (IBM MQ for z/OS only)

You might want to take advantage of the increased availability, scalability, and workload balancing that are possible when you use shared queues with queue-sharing groups. See *Shared queues and queue-sharing groups* for more information.

You might also want to estimate the average and peak message flows and consider using queue-sharing group clusters to spread the workload. See *Shared queues and queue-sharing groups* for more information.

Using queue manager clusters

You might want to take advantage of the simplified system administration, and increased availability, scalability, and workload balancing that are possible when you use clusters.

Types of messages

You might want to use datagrams for simple messages, but request messages (for which you expect replies) for other situations. You might want to assign different priorities to some of your messages. For more information about designing messages, see “Design techniques for messages” on page 46.

Using publish/subscribe or point-to-point messaging

Using publish/subscribe messaging, a sending application sends the information that it wants to share in an IBM MQ message to a standard destination managed by IBM MQ publish?subscribe, and lets IBM MQ handle the distribution of that information. The target application does not have to know anything about the source of the information it receives, it just registers an interest in one or more topics and receives that information when it is available. For more information about publish/subscribe messaging, see Publish/subscribe messaging.

Using point-to-point messaging, a sending application sends a message to a specific queue, from where it knows a receiving application will retrieve it. A receiving application gets messages from a specific queue and acts on their contents. An application will often function both as a sender and a receiver, sending a query to another application and receiving a response.

Controlling your IBM MQ programs

You might want to start some programs automatically or make programs wait until a particular message arrives on a queue (using the IBM MQ *triggering* feature, see “Starting IBM MQ applications using triggers” on page 234). Alternatively, you might want to start another instance of an application when the messages on a queue are not getting processed fast enough (using the IBM MQ *instrumentation events* feature as described in Instrumentation events).

Running your application on an IBM MQ client

The full MQI is supported in the client environment, and almost any IBM MQ application written in a procedural language can be relinked to run on an IBM MQ MQI client. Link the application on the IBM MQ MQI client to the MQIC library, rather than to the MQI library.

► **z/OS** Get(signal) on z/OS is not supported.

Note: An application running on an IBM MQ client can connect to more than one queue manager concurrently, or use a queue manager name with an asterisk (*) on an MQCONN or MQCONNX call. Change the application if you want to link to the queue manager libraries instead of the client libraries, as this function will not be available.

See “Running applications in the IBM MQ MQI client environment” on page 294 for more information.

Application performance

Design decisions can impact your application performance, for suggestions for enhancing performance of IBM MQ applications, see “Application design and performance considerations” on page 48 ► **IBM i** and “Design and performance considerations for IBM i applications” on page 51.

Advanced IBM MQ techniques

For more advanced applications you might want to use some advanced IBM MQ techniques such as correlating replies, and generating and sending IBM MQ context information. For more information, see “Design techniques for advanced applications” on page 50.

Securing your data and maintaining its integrity

You can use the context information that is passed with a message to test that the message has been sent from an acceptable source. You can use the syncpointing facilities provided by IBM MQ or your operating system to ensure that your data remains consistent with other resources (see “Committing and backing out units of work” on page 222 for further details). You can use the *persistence* feature of IBM MQ messages to assure the delivery of important messages.

Testing IBM MQ applications

The application development environment for IBM MQ programs is no different from that for any other application, so you can use the same development tools as well as the IBM MQ trace facilities.

z/OS When testing CICS applications with IBM MQ for z/OS, you can use the CICS Execution Diagnostic Facility (CEDF). CEDF traps the entry and exit of every MQI call as well as calls to all CICS services. Also, in the CICS environment, you can write an API-crossing exit program to provide diagnostic information before and after every MQI call. For information about how to do this, see “Using and writing applications on IBM MQ for z/OS” on page 258.

IBM i When testing IBM i applications, you can use the standard Debugger. To start this, use the STRDBG command.

Handling exceptions and errors

You need to consider how to process messages that cannot be delivered, and how to resolve error situations that are reported to you by the queue manager. For some reports, you must set report options on MQPUT.

Related concepts:

z/OS “Design and performance considerations for z/OS applications” on page 53
Application design is one of the most important factors affecting performance. Use this topic to understand some of the design factors involved in performance.

“Developing applications for IBM MQ” on page 1

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM MQ support applications written in procedural languages, and object oriented languages and frameworks.

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

“Writing a procedural application for queuing” on page 74

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

“Writing client procedural applications” on page 285

What you need to know to write client applications on IBM MQ using a procedural language.

“Developing .NET applications” on page 630

IBM MQ classes for .NET allow a program written in the .NET programming framework to connect to IBM MQ as an IBM MQ MQI client or to connect directly to an IBM MQ server.

“Developing C++ applications” on page 796

IBM MQ provides C++ classes equivalent to IBM MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI.

“Using IBM MQ classes for JMS” on page 824

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the javax.jms package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

“Using IBM MQ classes for Java” on page 1021

Use IBM MQ in a Java environment. IBM MQ classes for Java allow a Java application to connect to IBM MQ as an IBM MQ client, or connect directly to an IBM MQ queue manager.

“Using the Component Object Model Interface (IBM MQ Automation Classes for ActiveX)” on page 681

The IBM MQ Automation Classes for ActiveX (MQAX) are ActiveX components that provide classes that you can use in your application to access IBM MQ.

Related information:

IBM MQ technical overview

Choosing to use IBM MQ classes for Java or IBM MQ classes for JMS

A Java application can use either IBM MQ classes for Java or IBM MQ classes for JMS to access IBM MQ resources. Each approach has its advantages.

Note: IBM will make no further enhancements to the IBM MQ classes for Java and they are functionally stabilized at the level shipped in IBM MQ Version 8.0.

IBM MQ classes for Java encapsulates the Message Queue Interface (MQI), the native IBM MQ API, and uses the same object model as other object-oriented interfaces, whereas IBM MQ classes for Java Message Service implements Oracle's Java Message Service (JMS) interfaces.

If you are familiar with IBM MQ in environments other than Java, using either procedural or object-oriented languages, you can transfer your existing knowledge to the Java environment by using IBM MQ classes for Java. You can also exploit the full range of features of IBM MQ, not all of which are available in IBM MQ classes for JMS.

If you are not familiar with IBM MQ, or already have JMS experience, you might find it easier to use the familiar JMS API to access IBM MQ resources, by using IBM MQ classes for JMS. JMS is also an integral part of the Java Platform, Enterprise Edition (Java EE) platform. Java EE applications can use message-driven beans (MDBs) to process messages asynchronously. JMS is also the standard mechanism for Java EE to interact with asynchronous messaging systems such as IBM MQ. Every application server that is Java EE compliant must include a JMS provider, therefore you can use JMS to communicate between different application servers or you can port an application from one JMS provider to another without any change to the application.

“Using IBM MQ classes for Java” on page 1021

Use IBM MQ in a Java environment. IBM MQ classes for Java allow a Java application to connect to IBM MQ as an IBM MQ client, or connect directly to an IBM MQ queue manager.

“Using IBM MQ classes for JMS” on page 824

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the `javax.jms` package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

Scenarios: WebSphere Application Server with IBM MQ

Scenarios: WebSphere Application Server Liberty profile with IBM MQ

Design techniques for messages

Consider the aspects given in this information to help you design messages.

You create a message when you use an MQI call to put the message on a queue. As input to the call, you supply some control information in a *message descriptor* (MQMD) and the data that you want to send to another program. But at the design stage, you need to consider the following, because they affect the way that you create your messages:

Type of message to use

Are you designing a simple application in which you can send a message, then take no further action? Or are you asking for a reply to a question? If you are asking a question, you might include in the message descriptor the name of the queue on which you want to receive the reply.

Do you want your request and reply messages to be synchronous? This implies that you set a timeout period for the reply to answer your request, and if you do not receive the reply within that period, it is treated as an error.

Or would you prefer to work asynchronously, so that your processes do not have to depend upon the occurrence of specific events, such as common timing signals?

Another consideration is whether you have all your messages inside a unit of work.


Assigning different priorities to messages

You can assign a priority value to each message, and define the queue so that it maintains its messages in order of their priority. If you do this, when another program retrieves a message from the queue, it always gets the message with the highest priority. If the queue does not maintain its messages in priority order, a program that retrieves messages from the queue will retrieve them in the order in which they were added to the queue.

Programs can also select a message using the identifier that the queue manager assigned when the message was put on the queue. Alternatively, you can generate your own identifiers for each of your messages.

Effect of restarting queue manager on messages


The queue manager preserves all persistent messages, recovering them when necessary from the IBM MQ log files, when it is restarted. Nonpersistent messages and temporary dynamic queues are not preserved. Any messages that you do not want discarded must be defined as persistent when they are created. When writing an application for IBM MQ for Windows or IBM MQ on UNIX and Linux systems, make sure that you know how your system has been set up in respect of log file allocation to reduce the risk of designing an application that will run to the log file limits.


 Because messages on shared queues (only available on IBM MQ for z/OS) are held in the coupling facility (CF), nonpersistent messages are preserved across restarts of a queue manager as long as the CF remains available. If the CF fails, nonpersistent messages are lost.

Giving information about yourself to the recipient of messages

Usually, the queue manager sets the user ID, but suitably authorized applications can also set this field, so that you can include your own user ID and other information that the receiving program can use for accounting or security purposes.

Amount of receiving queues

 If a message might need to be put on several queues, you can publish to a topic or a distribution list.

 If a message might need to be put on several queues, you can publish to a topic.

Selectors and message properties

Messages can have metadata associated with them alongside the main message payload. These message properties can be useful in supplying additional data.

There are two aspects to this additional data that it is important to know about:

- The properties are not subject to Advanced Message Security (AMS) protection. If you want to use AMS to protect your data, then put it in the payload and not the message properties.
- The properties can be used to perform the selection of messages.

It is important to note that using selectors breaks the standard message convention of first in first out. As the queue manager is optimized for this workload, providing complex selectors is not advised for performance reasons. The queue manager does not store indexes of the message properties, therefore searching for a message must be a linear search. The deeper the queue, the more complex the selector, and the lower probability that the selector matching a message will adversely affect performance.

If complex selection is required, it is suggested to filter the messages by using any application or processing engine, such as IBM Integration Bus, to different destinations. Alternatively, the use of a topic hierarchy might be useful.

Note: IBM MQ classes for Java do not support the use of selectors, if you do wish to use selectors these should be done via the JMS API.

Application design and performance considerations

There are a number of ways in which poor program design can affect performance. These can be difficult to detect because the program can appear to perform well itself, but affect the performance of other tasks. Several problems specific to programs making IBM MQ calls are explained in this topic.

Here are a few ideas to help you to design efficient applications:


- Design your application so that processing goes on in parallel with a user's thinking time:
 - Display a panel and allow the user to start typing while the application is still initializing.
 - Get the data that you need in parallel from different servers.
- Keep connections and queues open if you are going to reuse them instead of repeatedly opening and closing, connecting, and disconnecting.
- However, a server application that is putting only one message should use MQPUT1.
- Queue managers are optimized for messages between 4 KB and 100 KB in size. Very large messages are inefficient; it is probably better to send 100 messages of 1 MB each than a single 100 MB message. Very small messages are also inefficient. The queue manager does the same amount of work for a single-byte message as for a 4 KB message.
- Keep your messages within a unit of work so that they can be committed or backed out simultaneously.
- Use the nonpersistent option for messages that do not need to be recoverable.
- If you need to send a message to a number of target queues, consider using a distribution list.

Effect of message length

The amount of data in a message can affect the performance of the application that processes the message. To achieve the best performance from your application, send only the essential data in a message. For example, in a request to debit a bank account, the only information that might need to be passed from the client to the server application is the account number and the amount of the debit.

Effect of message persistence

Persistent messages are usually logged. Logging messages reduces the performance of your application, so use persistent messages for essential data only. If the data in a message can be discarded if the queue manager stops or fails, use a nonpersistent message.

 MQPUT and MQGET operations for persistent messages will block when there is insufficient recovery log space to record the operations. Such a condition is indicated in the queue manager job log by messages CSQJ110E and CSQJ111A. Ensure monitoring processes are in place so that such conditions are managed and avoided.

Searching for a particular message

The MQGET call usually retrieves the first message from a queue. If you use the message and correlation identifiers (*MsgId* and *CorrelId*) in the message descriptor to specify a particular message, the queue manager has to search the queue until it finds that message. Using the MQGET call in this way affects the performance of your application.

Queues that contain messages of different lengths

If your application cannot use messages of a fixed length, grow and shrink the buffers dynamically to suit the typical message size. If the application issues an MQGET call that fails because the buffer is too small, the size of the message data is returned. Add code to your application so that the buffer is resized accordingly and the MQGET call is reissued.

Note: If you do not set the **MaxMsgLength** attribute explicitly, it defaults to 4 MB, which might be very inefficient if this is used to influence the application buffer size.

Frequency of sync points

Programs that issue very large numbers of MQPUT or MQGET calls within sync point, without committing them, can cause performance problems. Affected queues can fill up with messages that are currently inaccessible, while other tasks might be waiting to get these messages. This has implications in terms of storage, and in terms of threads that are tied up with tasks that are attempting to get messages.

Use of the MQPUT1 call

Use the MQPUT1 call only if you have a single message to put on a queue. If you want to put more than one message, use the MQOPEN call, followed by a series of MQPUT calls and a single MQCLOSE call.

Number of threads in use

Windows For IBM MQ for Windows, an application might require a large number of threads. Each queue manager process is allocated a maximum allowable number of application threads.

Applications might use too many threads. Consider whether the application takes into account this possibility and that it takes actions either to stop or to report this type of occurrence.

Put persistent messages under syncpoint

Persistent messages should be put and got under syncpoint. This is because when getting a persistent message outside of syncpoint, if the get fails, there is no way for the application to know whether the message has been got from the queue or not, and whether, if the message has been got, then it has also been lost. When getting persistent messages under syncpoint, if anything fails, the transaction is rolled back and the persistent message is not lost because it is still on the queue. Similarly, when putting persistent messages, put them under syncpoint. Another reason for putting and getting persistent messages under syncpoint is that the persistent message code in IBM MQ is heavily optimized for syncpoint. So putting and getting persistent messages under syncpoint is faster than putting and getting persistent messages outside of syncpoint.

However, it is faster to put and get non-persistent messages outside of syncpoint because the nonpersistent code in IBM MQ is optimized for being outside of syncpoint. Putting and getting persistent messages go at disk speeds because the persistent message is persisted to disk. However, putting and getting non-persistent messages go at CPU speeds because there is no disk write involved, not even when using syncpoint.

If an application is getting messages and does not know in advance whether they are persistent or not, the GMO option MQGMO_SYNCPOINT_IF_PERSISTENT can be used.


Design techniques for advanced applications

When designing more advanced applications, there are some techniques that you might want to consider such as waiting for messages, correlating replies, setting and using context information, starting applications automatically, generating reports and removing message affinities when using clustering.

For a simple IBM MQ application, you need to decide which IBM MQ objects to use in your application, and which types of message you want to use. For a more advanced application, you might want to use some of the techniques introduced in the following sections.

Waiting for messages

A program that is serving a queue can await messages by:

- Waiting until either a message arrives, or a specified time interval expires (see “Waiting for messages” on page 159).
-  On IBM MQ for z/OS only, setting a signal so that the program is informed when a message arrives. For more information, see “Signaling” on page 160.
- Establishing a callback exit to be driven when a message arrives; see “Asynchronous consumption of IBM MQ messages” on page 34.
- Making periodic calls on the queue to see whether a message has arrived (*polling*). This is not typically advisable because it can have performance implications.

Correlating replies

In IBM MQ applications, when a program receives a message that requests it to do some work, the program typically sends one or more reply messages to the requester.

To help the requester to associate these replies with its original request, an application can set a *correlation identifier* field in the descriptor of each message. Programs then copy the message identifier of the request message into the correlation identifier field of their reply messages.

Setting and using context information

Context information is used for associating messages with the user who generated them, and for identifying the application that generated the message. Such information is useful for security, accounting, auditing, and problem determination.

When you create a message, you can specify an option that requests that the queue manager associates default context information with your message.

For more information about using and setting context information, see “Message context” on page 40.

Starting IBM MQ programs automatically

Use IBM MQ *triggering* to start a program automatically when messages arrive on a queue.

You can set trigger conditions on a queue so that a program starts to process that queue:

- Every time that a message arrives on the queue
- When the first message arrives on the queue
- When the number of messages on the queue reaches a predefined number

For more information about triggering, see “Starting IBM MQ applications using triggers” on page 234. Triggering is just one way of starting a program automatically. For example, you can start a program automatically on a timer using non-IBM MQ facilities.

On Multiplatforms, IBM MQ can define service objects to start IBM MQ programs when the queue manager starts; see Service objects.

Generating IBM MQ reports

You can request the following reports within an application:

- Exception reports
- Expiry reports
- Confirm-on-arrival (COA) reports
- Confirm-on-delivery (COD) reports
- Positive action notification (PAN) reports
- Negative action notification (NAN) reports

These are described in “Report messages” on page 11.

Clusters and message affinities

Before starting to use clusters with multiple definitions for the same queue, examine your applications to see whether there are any that require an exchange of related messages.

Within a cluster, a message can be routed to any queue manager that hosts an instance of the appropriate queue. Therefore, the logic of applications with message affinities can be upset.

For example, you might have two applications that rely on a series of messages flowing between them in the form of questions and answers. It might be important that all the questions are sent to the same queue manager and that all the answers are sent back to the other queue manager. In this situation, it is important that the workload management routine does not send the messages to any queue manager that just happens to host an instance of the appropriate queue.

Where possible, remove the affinities. Removing message affinities improves the availability and scalability of applications.

For more information, see Handling message affinities.

Design and performance considerations for IBM i applications

Use this information to understand how application design, threads, and storage, can affect performance.

This information is split into two sections:

- “Application design considerations”
- “Specific performance problems” on page 52

Application design considerations

There are a number of ways in which poor program design can affect performance. These problems can be difficult to detect because the program can appear to perform well, while affecting the performance of other tasks. Several problems specific to programs making IBM MQ for IBM i calls are explained in the following sections.

For more information about application design, see “Design considerations for IBM MQ applications” on page 43.

Effect of message length

Although IBM MQ for IBM i allows messages to hold up to 100 MB of data, the amount of data in a message affects the performance of the application that processes the message. To achieve the

best performance from your application, send only the essential data in a message; for example, in a request to debit a bank account, the only information that might need to be passed from the client to the server application is the account number and the amount of the debit.

Effect of message persistence

Persistent messages are journaled. Journaling messages reduces the performance of your application, so use persistent messages for essential data only. If the data in a message can be discarded if the queue manager stops or fails, use a nonpersistent message.

Searching for a particular message

The MQGET call usually retrieves the first message from a queue. If you use the message and correlation identifiers (*MsgId* and *CorrelId*) in the message descriptor to specify a particular message, the queue manager must search the queue until it finds that message. The use of the MQGET call in this way affects the performance of your application.

Queues that contain messages of different lengths

If the messages on a queue are of different lengths, to determine the size of a message, your application can use the MQGET call with the *BufferLength* field set to zero so that, even though the call fails, it returns the size of the message data. The application can then repeat the call, specifying the identifier of the message it measured in its first call and a buffer of the correct size. However, if there are other applications serving the same queue, you might find that the performance of your application is reduced because its second MQGET call spends time searching for a message that another application has retrieved in the time between your two calls.

If your application cannot use messages of a fixed length, another solution to this problem is to use the MQINQ call to find the maximum size of messages that the queue can accept, then use this value in your MQGET call. The maximum size of messages for a queue is stored in the **MaxMsgLen** attribute of the queue. This method might use large amounts of storage, however, because the value of this queue attribute can be the maximum allowed by IBM MQ for IBM i, which might be greater than 2 GB.

Frequency of sync points

Programs that issue numerous MQPUT calls within sync point, without committing them, can cause performance problems. Affected queues can fill up with messages that are currently unusable, while other tasks might be waiting to get these messages. This problem has implications in terms of storage, and in terms of threads tied up with tasks that are attempting to get messages.

Use of the MQPUT1 call

Use the MQPUT1 call only if you have a single message to put on a queue. If you want to put more than one message, use the MQOPEN call, followed by a series of MQPUT calls and a single MQCLOSE call.

Number of threads in use

An application might require many threads. Each queue manager process is allocated a maximum allowable number of threads. If some applications are troublesome, it might be due to their design using too many threads. Consider whether the application takes into account this possibility and that it takes actions either to stop or to report this type of occurrence. The maximum number of threads that IBM i allows is 4,095. However, the default is 64. IBM MQ makes available up to 63 threads to its processes.

Specific performance problems

This section explains the problems of storage and poor performance.

Storage problems

If you receive the system message CPF0907. Serious storage condition may exist it is possible that you are filling up the space associated with the IBM MQ for IBM i queue managers.

Is your application or IBM MQ for IBM i running slowly?

If your application is running slowly, it might indicate that it is in a loop, or waiting for a resource that is not available. This slow running might also be caused by a performance problem. Perhaps it is because your system is operating near the limits of its capacity. This type of problem is probably worst at peak system load times, typically at mid-morning and mid-afternoon. (If your network extends across more than one time zone, peak system load might seem to you to occur at some other time.)

If you find that performance degradation is not dependent on system loading, but happens sometimes when the system is lightly loaded, a poorly designed application program is probably to blame. This problem might manifest itself as a problem that only occurs when certain queues are accessed.

QTOTJOB and QADLTOTJ are system values worth investigating.

The following symptoms might indicate that IBM MQ for IBM i is running slowly:

- If your system is slow to respond to MQSC commands.
- If repeated displays of the queue depth indicate that the queue is being processed slowly for an application with which you would expect a large amount of queue activity.
- Is IBM MQ trace running?

Linux on POWER Systems - Little Endian applications



As Linux on POWER[®] Systems - Little Endian supports 64-bit applications only, there is no support provided in IBM MQ for 32-bit applications.

Related concepts:

“Design considerations for IBM MQ applications” on page 43

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

Design and performance considerations for z/OS applications



Application design is one of the most important factors affecting performance. Use this topic to understand some of the design factors involved in performance.

There are a number of ways in which poor program design can affect performance. These problems can be difficult to detect because the program can appear to perform well, while affecting the performance of other tasks. Several problems specific to programs making MQI calls are demonstrated in the following sections.

For more information about application design, see “Design considerations for IBM MQ applications” on page 43.

Effect of message length

Although IBM MQ for z/OS allows messages to hold up to 100 MB of data, the amount of data in a message affects the performance of the application that processes the message. To achieve the best performance from your application, send only the essential data in a message. For example, in a request to debit a bank account, the only information that might need to be passed from the client to the server application is the account number and the amount to debit.

Effect of message persistence

Persistent messages are logged. Logging messages reduces the performance of your application, so use persistent messages for essential data only. If the data in a message can be discarded if the queue manager stops or fails, use a nonpersistent message.

Data for persistent messages is written to log buffers. These buffers are written to the log data sets when:

- A commit occurs
- A message is got or put out of syncpoint
- WRTHRSR buffers are filled

Processing many messages in one unit of work can cause less input/output than if the messages were processed one for each unit of work, or out of syncpoint.

Searching for a particular message

The MQGET call typically retrieves the first message from a queue. If you use the message and correlation identifiers (**MsgId** and **CorrelId**) in the message descriptor to specify a particular message, the queue manager searches the queue until it finds that message. Using MQGET in this way affects the performance of your application because, to find a particular message, IBM MQ might have to scan the entire queue.

You can use the **IndexType** queue attribute to specify that you want the queue manager to maintain an index that can be used to increase the speed of MQGET operations on the queue. However, there is a small performance reduction for maintaining an index, so only generate one if you need to use it. You can choose to build an index of message identifiers or of correlation identifiers, or you can choose not to build an index for queues where messages are retrieved sequentially. Try to have many different key values, not lots with the same value. For example Balance1, Balance2, and Balance3, not three with Balance. For shared queues, you must have the correct **IndexType**. For details of the **IndexType** queue attribute, see **IndexType**.

To avoid affecting queue manager restart time by using indexed queues, use the QINDXBLD(NOWAIT) parameter in the CSQ6SYSP macro. This allows the queue manager restart to complete without waiting for queue index building to complete.

For a full description of the **IndexType** attribute, and other object attributes see **Attributes of objects**.

Queues that contain messages of different lengths

Get a message, using a buffer size matching the expected size of the message. If you receive the return code indicating that the message is too long, get a bigger buffer. When the get fails in this way, the data length returned is the size of the unconverted message data. If you specify MQGMO_CONVERT on the MQGET call, and the data expands during conversion, it still might not fit in the buffer, in which case you need to further increase the size of the buffer.

If you issue the MQGET with a buffer length of zero, it returns the size of the message and the application can then get a buffer of this size and reissue the get. If you have multiple applications processing the queue, another application might have already processed the message when the original application reissued the get. If you occasionally have large messages, you might need to get a large buffer just for these messages, and release it after the message has been processed. This should help reduce virtual storage problems if all applications have large buffers.

If your application cannot use messages of a fixed length, another solution to this problem is to use the MQINQ call to find the maximum size of messages that the queue can accept, then use this value in your MQGET call. The maximum size of messages for a queue is stored in the **MaxMsgL** attribute of the queue. This method could use large amounts of storage, however, because the value of **MaxMsgL** could be as high as 100 MB, the maximum allowed by IBM MQ for z/OS.

Note: You can lower the **MaxMsgL** parameter after large messages have been put to the queue. For example you can put a 100 MB message, then set **MaxMsgL** to 50 bytes. This means that it is still possible to get bigger messages than the application expected.

Frequency of syncpoints

Programs that issue many MQPUT calls within syncpoint, without committing them, can cause performance problems. Affected queues can fill up with messages that are currently unusable, while other tasks might be waiting to get these messages. This has implications in terms of storage, and in terms of threads tied up with tasks that are attempting to get messages.

As a rule if you have multiple applications processing a queue you typically get the best performance when you have either

- 100 short messages (less than 1 KB), or
- One message for larger messages (100 KB)

for each syncpoint. If there is only one application processing the queue, you must have more messages for each unit of work.

You can limit the number of messages that a task can get or put within a single unit of recovery with the **MAXUMSGS** queue manager attribute. For information about this attribute, see the **ALTER QMGR** command in Script (MQSC) Commands.

Advantages of the MQPUT1 call

Use the MQPUT1 call only if you have a single message to put on a queue. If you want to put more than one message, use the MQOPEN call, followed by a series of MQPUT calls and a single MQCLOSE call.

How many messages can a queue manager contain

Local Queues

The number of local messages a queue manager can hold is basically the size of the page sets. You can have up to 100 page sets (though it is recommended page set 0 and page set 1 are for system related objects and queues). You can use a page set with extended format and increase the capacity of a page set.

Shared Queues

The capacity for shared queues depends on the size of the coupling facility (CF). IBM MQ uses CF list structures where fundamental storage units are entries and elements. Each message is stored as 1 entry and multiple elements containing the associated MQMD and other message data. The number of elements consumed by a single message depends on the size of the message and, for CFLEVEL(5), the offload rules in effect at MQPUT time. Fewer elements are needed when message data is offloaded to either Db2 or SMDS. Message data access is slower when the message has been offloaded. See Performance Supportpac MP1H for further comparison of performance and CPU overhead associated with message offload.

What affects performance

Performance can mean how fast messages can be processed, and it can also mean how much CPU is needed per message.

What affects how fast messages can be processed

For persistent messages the biggest impact is the speed of the log data sets. The speed of the log data sets depends on the DASD they are on. Therefore care should be taken to put log data set on low used volumes to reduce contention. Striping the MQ logs improves the log performance when there are multiple pages written per I/O. Z High Performance Fibre connection (zHPF) also

has a significant performance to I/O response time when the I/O subsystem is busy. See <http://www.ibm.com/support/docview.wss?uid=swg27015698>.

When there is a request to get and put a message, access to the queue is locked during the request to preserve integrity of the queue. For planning purposes consider the queue locked for the whole request. So if the time for a put is 100 microseconds, and you have more than 10,000 requests a second you might experience delays. You might achieve better than this in practice, but it is a good general rule. You can use different queues to improve performance.

Possible reasons for this can be:

- use a common reply queue which every CICS transaction uses
- each CICS transaction is given a unique reply to queue
- a reply to a queue for CICS region and all transactions in the CICS region use this queue.

The answer depends on the number of requests a second, and the response time of the requests.

If messages have to be read from a page set, they will be slower compared to when the messages are in the buffer pool. If you have more messages than fit into a buffer pool, then they will spill to disk. So you need to ensure that the buffer pool is big enough for your short lived messages. If you have messages that you process many hours later, these are likely to spill to disk, so you should expect a get for these messages to be slower than if they were in the buffer pool.

For a shared queue, the speed of the messages depends on the speed of the Coupling Facility. A CF within the physical processor is likely to be faster than an external CF. The CF response time depends on how busy the CF is. For example on the Hursley systems, when the CF was 17% busy the response time was 14 microseconds. When the CF was 95% busy the response time was 45 microseconds.

If your MQ requests use a lot of CPU, this can affect how fast messages are processed. Because if the Logical Partition (LPAR) is constrained for CPU, applications will be delayed waiting for CPU.

How much CPU per message

In general bigger messages use more CPU, so avoid large (x MB) messages if possible.

When getting specific messages from queues, the queue should be indexed so the queue manager can go directly to the message (and so avoids potentially an entire scan of the queue). If the queue is not indexed then the queue is scanned from the beginning looking for the message. If there are 1000 messages on the queue, it may have to scan all 1000 messages. The result is a lot of unnecessary CPU usage.

Channels using TLS have an additional cost due to the encryption of the message.

In MQ V7 you can select messages by a selector string in addition to the **CORRELID** or **MSGID**. Every message has to be looked in, so if there are many messages on the queue this is expensive.

It is more efficient for an application to do OPEN PUT PUT CLOSE than PUT1 PUT1.

Triggering in CICS

When the message arrival rate of messages for a triggered queue is low, it is efficient to use trigger first. When the message arrival rate is more than 10 messages a second, it is more efficient to trigger the first transaction, then have the transaction process a message and get the next message, and so on. If a message has not arrived in a short period (say between 0.1 and 1 second) the transaction ends. At high throughput you might need multiple transactions running to process the messages and to prevent a build up of messages. For every trigger message produced, this requires a put and a get of a trigger message, which in effect doubles the cost of the message.

How many connections or concurrent users are supported

Each connection uses virtual storage within the queue manager, so the more concurrent users the more storage used. If you need a very large buffer pool and large number of users, then you might be constrained for virtual storage, and you might need to reduce the size of your buffer pools.

If security is being used, the queue manager caches information within the queue manager for a long period. The amount of virtual storage that is used within the queue manager is affected.

The **CHINIT** can support up to about 10,000 connections. This is limited by virtual storage. If a connection uses more storage, for example using by TLS, the storage per connection increases, which therefore means the **CHINIT** can support less connections. If you are processing large messages, these will require more storage for buffers in the **CHINIT**, so the **CHINIT** can support less messages.

Connections to a remote queue manager are more efficient than client connections. For example, every MQ client requests requires two network flows (one for the request and one for the response). With a channel to a remote queue manager, there may be 50 sends over the network before a response comes back. If you are considering a large client network, it may be more efficient to use a concentrator queue manager on a distributed box, and have one channel coming in and out of the concentrator.

Other things affecting performance


Log data set should be at least 1000 cylinders in size. If the logs are smaller than this, checkpoint activity may be too frequent. On a busy system a checkpoint typically should be every 15 minutes or longer, at very high throughputs it may less than this. When a checkpoint occurs the buffer pools are scanned and 'old' messages and changed pages are written to disk. If checkpoints are too frequent, this can impact performance. The value of LOGLOAD can also affect checkpoint frequency. If the queue manager abnormally ends, then at restart it may have to read back to 3 checkpoints. The best checkpoint interval is a balance between the activity when a checkpoint is taken, and the amount of log data that may need to be read when the queue manager restarts.

There is a significant overhead incurred when starting a channel. It is usually better to start a channel and leave it connected, rather than frequent starts and stops of the channel.

Related information:

 [MP1H: IBM MQ for z/OS Version 9.0 Performance Report](#)

IMS and IMS bridge applications on IBM MQ for z/OS

 z/OS

This information helps you to write IMS applications using IBM MQ.

- To use syncpoints and MQI calls in IMS applications, see “Writing IMS applications using IBM MQ” on page 58.
- To write applications that use the IBM MQ - IMS bridge, see “Writing IMS bridge applications” on page 63.

Use the following links to find out more about IMS and IMS bridge applications on IBM MQ for z/OS:

- “Writing IMS applications using IBM MQ” on page 58
- “Writing IMS bridge applications” on page 63

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

“Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

Writing IMS applications using IBM MQ

There are further considerations when using IBM MQ in IMS applications. These include which MQ API calls can be used and the mechanism used for syncpoint.

Use the following links to find out more about writing IMS applications on IBM MQ for z/OS:

- “Syncpoints in IMS applications” on page 59
- “MQI calls in IMS applications” on page 59

Restrictions

There are restrictions on which IBM MQ API calls can be used by an application using the IMS adapter.

The following IBM MQ API calls are not supported within an application using the IMS adapter:

- MQCB
- MQCB_FUNCTION
- MQCTL

Related concepts:

“Writing IMS bridge applications” on page 63

This topic contains information about writing applications to use the IBM MQ - IMS bridge.

Syncpoints in IMS applications:

In an IMS application, you establish a syncpoint by using IMS calls such as GU (get unique) to the IOPCB and CHKP (checkpoint).

To back out all changes since the previous checkpoint, you can use the IMS ROLB (rollback) call. For more information, see the following documentation:

- IMS Version 13 Application Programming APG SC19-3646
- IMS Version 13 Application Programming APIs APR SC19-3647

The queue manager is a participant in a two-phase commit protocol; the IMS syncpoint manager is the coordinator.

All open handles are closed by the IMS adapter at a syncpoint (except in a batch or non-message driven BMP environment). This is because a different user could initiate the next unit of work and IBM MQ security checking is performed when the MQCONN, MQCONNX, and MQOPEN calls are made, not when the MQPUT or MQGET calls are made.

However, in a Wait-for-Input (WFI) or pseudo Wait-for-Input (PWFI) environment IMS does not notify IBM MQ to close the handles until either the next message arrives or a QC status code is returned to the application. If the application is waiting in the IMS region and any of these handles belong to triggered queues, triggering will not occur because the queues are open. For this reason, applications running in a WFI or PWFI environment should explicitly MQCLOSE the queue handles before doing the GU to the IOPCB for the next message.

If an IMS application (either a BMP or an MPP) issues the MQDISC call, open queues are closed but no implicit syncpoint is taken. If the application ends normally, any open queues are closed and an implicit commit occurs. If the application ends abnormally, any open queues are closed and an implicit backout occurs.

MQI calls in IMS applications:

Use this information to learn about the use of MQI calls on Server applications and Enquiry applications.

This section covers the use of MQI calls in the following types of IMS applications:

- “Server applications”
- “Inquiry applications” on page 62

Server applications

Here is an outline of the MQI server application model:

```
Initialize/Connect
.
Open queue for input shared
.
Get message from IBM MQ queue
.
Do while Get does not fail
.
If expected message received
Process the message
Else
Process unexpected message
```

```

End if
.
Commit
.
Get next message from IBM MQ queue
.
End do
.
Close queue/Disconnect
.
END

```

Sample program CSQ4ICB3 shows the implementation, in C/370™, of a BMP using this model. The program establishes communication with IMS first, and then with IBM MQ:

```

main()
----
Call InitIMS
If IMS initialization successful
Call InitMQM
If IBM MQ initialization successful
Call ProcessRequests
Call EndMQM
End-if
End-if

Return

```

The IMS initialization determines whether the program has been called as a message-driven or a batch-oriented BMP and controls IBM MQ queue manager connection and queue handles accordingly:

```

InitIMS
-----
Get the IO, Alternate and Database PCBs
Set MessageOriented to true

Call ctdli to handle status codes rather than abend
If call is successful (status code is zero)
While status code is zero
Call ctdli to get next message from IMS message queue
If message received
Do nothing
Else if no IOPBC
Set MessageOriented to false
Initialize error message
Build 'Started as batch oriented BMP' message
Call ReportCallError to output the message
End-if
Else if response is not 'no message available'
Initialize error message
Build 'GU failed' message
Call ReportCallError to output the message
Set return code to error
End-if
End-if
End-while
Else
Initialize error message
Build 'INIT failed' message
Call ReportCallError to output the message
Set return code to error
End-if

Return to calling function

```

The IBM MQ initialization connects to the queue manager and opens the queues. In a message-driven BMP this is called after each IMS syncpoint is taken; in a batch-oriented BMP, this is called only during program startup:

```

InitMQM
-----
Connect to the queue manager
If connect is successful
Initialize variables for the open call
Open the request queue
If open is not successful
Initialize error message
Build 'open failed' message
Call ReportCallError to output the message
Set return code to error
End-if
Else
Initialize error message
Build 'connect failed' message
Call ReportCallError to output the message
Set return code to error
End-if

Return to calling function

```

The implementation of the server model in an MPP is influenced by the fact that the MPP processes a single unit of work per invocation. This is because, when a syncpoint (GU) is taken, the connection and queue handles are closed and the next IMS message is delivered. This limitation can be partially overcome by one of the following:

- **Processing many messages within a single unit-of-work**

This involves:

- Reading a message
- Processing the required updates
- Putting the reply

in a loop until all messages have been processed or until a set maximum number of messages has been processed, at which time a syncpoint is taken.

Only certain types of application (for example, a simple database update or inquiry) can be approached in this way. Although the MQI reply messages can be put with the authority of the originator of the MQI message being handled, the security implications of any IMS resource updates need to be addressed carefully.

- **Processing one message per invocation of the MPP and ensuring multiple scheduling of the MPP to process all available messages.**

Use the IBM MQ IMS trigger monitor program (CSQQTRMN) to schedule the MPP transaction when there are messages on the IBM MQ queue and no applications serving it.

If trigger monitor starts the MPP, the queue manager name and queue name are passed to the program, as shown in the following COBOL code extract:

```

* Data definition extract
01 WS-INPUT-MSG.
05 IN-LL1          PIC S9(3) COMP.
05 IN-ZZ1          PIC S9(3) COMP.
05 WS-STRINGPARM  PIC X(1000).
01 TRIGGER-MESSAGE.
COPY CMQTM2L.
*
* Code extract
GU-IOPCB SECTION.
MOVE SPACES TO WS-STRINGPARM.
CALL 'CBLTDLI' USING GU,
IOPCB,

```

```

WS-INPUT-MSG.
IF IOPCB-STATUS = SPACES
MOVE WS-STRINGPARM TO MQTMC.
* ELSE handle error
*
* Now use the queue manager and queue names passed
DISPLAY 'MQTMC-QMGRNAME      ='
MQTMC-QMGRNAME OF MQTMC '='.
DISPLAY 'MQTMC-QNAME        ='
MQTMC-QNAME OF MQTMC '='.

```

The server model, which is expected to be a long running task, is better supported in a batch processing region, although the BMP cannot be triggered using CSQQTRMN.

Inquiry applications

A typical IBM MQ application initiating an inquiry or update works as follows:

- Gather data from the user
- Put one or more IBM MQ messages
- Get the reply messages (you might have to wait for them)
- Provide a response to the user

Because messages put on to IBM MQ queues do not become available to other IBM MQ applications until they are committed, they must either be put out of syncpoint, or the IMS application must be split into two transactions.

If the inquiry involves putting a single message, you can use the *no syncpoint* option; however, if the inquiry is more complex, or resource updates are involved, you might get consistency problems if failure occurs and you do not use syncpointing.

To overcome this, you can split IMS MPP transactions using MQI calls using a program-to-program message switch; see *IMS/ESA Application Programming: Data Communication* for information about this. This allows an inquiry program to be implemented in an MPP:

```

Initialize first program/Connect
.
Open queue for output
.
Put inquiry to IBM MQ queue
.
Switch to second IBM MQ program, passing necessary data in save
pack area (this commits the put)
.
END
.
.
Initialize second program/Connect
.
Open queue for input shared
.
Get results of inquiry from IBM MQ queue
.
Return results to originator
.
END

```

Writing IMS bridge applications

This topic contains information about writing applications to use the IBM MQ - IMS bridge.

For information about the IBM MQ - IMS bridge, see *The IMS bridge*.

Use the following links to find out more about writing IMS bridge applications on IBM MQ for z/OS:

- “How the IMS bridge deals with messages”
- “Writing IMS transaction programs through IBM MQ” on page 283

Related concepts:

“Writing IMS applications using IBM MQ” on page 58

There are further considerations when using IBM MQ in IMS applications. These include which MQ API calls can be used and the mechanism used for syncpoint.

How the IMS bridge deals with messages:

When you use the IBM MQ - IMS bridge to send messages to an IMS application, you need to construct your messages in a special format.

You must also put your messages on IBM MQ queues that have been defined with a storage class that specifies the XCF group and member name of the target IMS system. These are known as MQ-IMS bridge queues, or simply **bridge** queues.

The IBM MQ-IMS bridge requires exclusive input access (MQOO_INPUT_EXCLUSIVE) to the bridge queue if it is defined with QSGDISP(QMGR), or if it is defined with QSGDISP(SHARED) together with the NOSHARE option.

A user does not need to sign on to IMS before sending messages to an IMS application. The user ID in the *UserIdentifier* field of the MQMD structure is used for security checking. The level of checking is determined when IBM MQ connects to IMS, and is described in Application access control for the IMS bridge. This enables a pseudo signon to be implemented.

The IBM MQ - IMS bridge accepts the following types of message:

- Messages containing IMS transaction data and an MQIIH structure (described in MQIIH):
MQIIH LLZZ<trancode><data>[LLZZ<data>][LLZZ<data>]

Note:

1. The square brackets, [], represent optional multi-segments.
 2. Set the *Format* field of the MQMD structure to MQFMT_IMS to use the MQIIH structure.
- Messages containing IMS transaction data but no MQIIH structure:
LLZZ<trancode><data> \
[LLZZ<data>][LLZZ<data>]

IBM MQ validates the message data to ensure that the sum of the LL bytes plus the length of the MQIIH (if it is present) is equal to the message length.

When the IBM MQ - IMS bridge gets messages from the bridge queues, it processes them as follows:

- If the message contains an MQIIH structure, the bridge verifies the MQIIH (see MQIIH), builds the OTMA headers, and sends the message to IMS. The transaction code is specified in the input message. If this is an LTERM, IMS replies with a DFS1288E message. If the transaction code represents a command, IMS executes the command; otherwise the message is queued in IMS for the transaction.
- If the message contains IMS transaction data, but no MQIIH structure, the IMS bridge makes the following assumptions:
 - The transaction code is in bytes 5 through 12 of the user data

- The transaction is in nonconversational mode
- The transaction is in commit mode 0 (commit-then-send)
- The *Format* in the MQMD is used as the *MFSMapName* (on input)
- The security mode is MQISS_CHECK

The reply message is also built without an MQIIH structure, taking the *Format* for the MQMD from the *MFSMapName* of the IMS output.

The IBM MQ - IMS bridge uses one or two Tpipes for each IBM MQ queue:

- A synchronized Tpipe is used for all messages using Commit mode 0 (COMMIT_THEN_SEND) (these show with SYN in the status field of the IMS /DIS TMEMBER client TPIPE xxxx command)
- A non-synchronized Tpipe is used for all messages using Commit mode 1 (SEND_THEN_COMMIT)

The Tpipes are created by IBM MQ when they are first used. A non-synchronized Tpipe exists until IMS is restarted. Synchronized Tpipes exist until IMS is cold started. You cannot delete these Tpipes yourself.

See the following topics for more information about how the IBM MQ - IMS bridge deals with messages:

- "Mapping IBM MQ messages to IMS transaction types"
- "If the message cannot be put to the IMS queue" on page 65
- "IMS bridge feedback codes" on page 65
- "The MQMD fields in messages from the IMS bridge" on page 66
- "The MQIIH fields in messages from the IMS bridge" on page 67
- "Reply messages from IMS" on page 68
- "Using alternate response PCBs in IMS transactions" on page 68
- "Sending unsolicited messages from IMS" on page 68
- "Message segmentation" on page 69
- "Data conversion" on page 69

Related concepts:

"Writing IMS transaction programs through IBM MQ" on page 283

The coding required to handle IMS transactions through IBM MQ depends on the message format required by the IMS transaction and the range of responses it can return. However, there are several points to consider when your application handles IMS screen formatting information.

Mapping IBM MQ messages to IMS transaction types:

A table describing the mapping of IBM MQ messages to IMS transaction types.

Table 4. Mapping IBM MQ messages to IMS transaction types

IBM MQ message type	Commit-then-send (mode 0) - uses synchronized IMS Tpipes	Send-then-commit (mode 1) - uses non-synchronized IMS Tpipes
Persistent IBM MQ messages	<ul style="list-style-type: none"> • Recoverable full function transactions • Unrecoverable transactions are rejected by IMS 	<ul style="list-style-type: none"> • Fastpath transactions • Conversational transactions • Full function transactions
Nonpersistent IBM MQ messages	<ul style="list-style-type: none"> • Unrecoverable full function transactions • Recoverable transactions are permitted with IMS V8 and APAR PQ61404 and all later versions of IMS 	<ul style="list-style-type: none"> • Fastpath transactions • Conversational transactions • Full function transactions

Note: IMS commands cannot use persistent IBM MQ messages with commit mode 0. See the *IMS/ESA Open Transaction Manager Access User's Guide* for more information.

If the message cannot be put to the IMS queue:

Learn about actions to take if the message cannot be put to the IMS queue.

If the message cannot be put to the IMS queue, the following action is taken by IBM MQ:

- If a message cannot be put to IMS because the message is invalid, the message is put to the dead-letter queue, and a message is sent to the system console.
- If the message is valid, but is rejected by IMS, IBM MQ sends an error message to the system console, the message includes the IMS sense code, and the IBM MQ message is put to the dead-letter queue. If the IMS sense code is 001A, IMS sends an IBM MQ message containing the reason for the failure to the reply-to queue.

Note: In the circumstances listed previously, if IBM MQ cannot put the message to the dead-letter queue for any reason, the message is returned to the originating IBM MQ queue. An error message is sent to the system console, and no further messages are sent from that queue.

To resend the messages, do *one* of the following:

- Stop and restart the Tpipes in IMS corresponding to the queue
 - Alter the queue to GET(DISABLED), and again to GET(ENABLED)
 - Stop and restart IMS or the OTMA
 - Stop and restart your IBM MQ subsystem
- If the message is rejected by IMS for anything other than a message error, the IBM MQ message is returned to the originating queue, IBM MQ stops processing the queue, and an error message is sent to the system console.

If an exception report message is required, the bridge puts it to the reply-to queue with the authority of the originator. If the message cannot be put to the queue, the report message is put to the dead-letter queue with the authority of the bridge. If it cannot be put to the DLQ, it is discarded.

IMS bridge feedback codes:

IMS sense codes are typically output in hexadecimal format in IBM MQ console messages such as CSQ2001I (for example, sense code 0x001F). IBM MQ feedback codes as seen in the dead-letter header of messages put to the dead-letter queue are decimal numbers.

The IMS bridge feedback codes are in the range 301 through 399, or 600 through 855 for NACK sense code 0x001A. They are mapped from the IMS-OTMA sense codes as follows:

1. The IMS-OTMA sense code is converted from a hexadecimal number to a decimal number.
2. 300 is added to the number resulting from the calculation in 1, giving the IBM MQ *Feedback* code.
3. The IMS-OTMA sense code 0x001A, decimal 26 is a special case. A *Feedback* code in the range 600-855 is generated.
 - a. The IMS-OTMA reason code is converted from a hexadecimal number to a decimal number.
 - b. 600 is added to the number resulting from the calculation in a, giving the IBM MQ *Feedback* code.

For information about IMS-OTMA sense codes, see the *IMS Messages and Codes*.

The MQMD fields in messages from the IMS bridge:

Learn about the MQMD fields in messages from the IMS bridge.

The MQMD of the originating message is carried by IMS in the User Data section of the OTMA headers. If the message originates in IMS, this is built by the IMS Destination Resolution Exit. The MQMD of a message received from IMS is built as follows:

StrucID

"MD "

Version

MQMD_VERSION_1

Report

MQRO_NONE

MsgType

MQMT_REPLY

Expiry If MQIIH_PASS_EXPIRATION is set in the Flags field of the MQIIH, this field contains the remaining expiry time, else it is set to MQEI_UNLIMITED

Feedback

MQFB_NONE

Encoding

MQENC.Native (the encoding of the z/OS system)

CodedCharSetId

MQCCSI_Q_MGR (the CodedCharSetID of the z/OS system)

Format

MQFMT_IMS if the MQMD.Format of the input message is MQFMT_IMS, otherwise IOPCB.MODNAME

Priority

MQMD.Priority of the input message

Persistence

Depends on commit mode: MQMD.Persistence of the input message if CM-1; persistence matches recoverability of the IMS message if CM-0

MsgId MQMD.MsgId if MQRO_PASS_MSG_ID, otherwise New MsgId (the default)

CorrelId

MQMD.CorrelId from the input message if MQRO_PASS_CORREL_ID, otherwise MQMD.MsgId from the input message (the default)

BackoutCount

0

ReplyToQ

Blanks

ReplyToQMGr

Blanks (set to local qmgr name by the queue manager during the MQPUT)

UserIdentifier

MQMD.UserIdentifier of the input message

AccountingToken

MQMD.AccountingToken of the input message

ApplIdentityData

MQMD.ApplIdentityData of the input message

PutApplType

MQAT_XCF if no error, otherwise MQAT_BRIDGE

PutApplName

<XCFgroupName><XCFmemberName> if no error, otherwise QMGR name

PutDate

Date when message was put

PutTime

Time when message was put

ApplOriginData

Blanks

The MQIIH fields in messages from the IMS bridge:

Learn about the MQIIH fields in messages from the IMS bridge.

The MQIIH of a message received from IMS is built as follows:

StrucId

"IIH "

Version

1

StrucLength

84

Encoding

MQENC_NATIVE

CodedCharSetId

MQCCSI_Q_MGR

Format

MQIIH.ReplyToFormat of the input message if MQIIH.ReplyToFormat is not blank, otherwise IOPCB.MODNAME

Flags 0

LTermOverride

LTERM name (Tpipe) from OTMA header

MFSMapName

Map name from OTMA header

ReplyToFormat

Blanks

Authenticator

MQIIH.Authenticator of the input message if the reply message is being put to an MQ-IMS bridge queue, otherwise blanks.

TranInstanceId

Conversation ID / Server Token from OTMA header if in conversation, otherwise nulls

TranState

"C" if in conversation, otherwise blank

CommitMode

Commit mode from OTMA header ("0" or "1")

SecurityScope

Blank

Reserved

Blank

Reply messages from IMS:

When an IMS transaction ISRTs to its IOPCB, the message is routed back to the originating LTERM or TPIPE.

These are seen in IBM MQ as reply messages. Reply messages from IMS are put onto the reply-to queue specified in the original message. If the message cannot be put onto the reply-to queue, it is put onto the dead-letter queue using the authority of the bridge. If the message cannot be put onto the dead-letter queue, a negative acknowledgment is sent to IMS to say that the message cannot be received. Responsibility for the message is then returned to IMS. If you are using commit mode 0, messages from that Tpipe are not sent to the bridge, and remain on the IMS queue; that is, no further messages are sent until restart. If you are using commit mode 1, other work can continue.

If the reply has an MQIIH structure, its format type is MQFMT_IMS; if not, its format type is specified by the IMS MOD name used when inserting the message.

Using alternate response PCBs in IMS transactions:

When an IMS transaction uses alternate response PCBs (ISRTs to the ALTPCB, or issues a CHNG call to a modifiable PCB), the pre-routing exit (DFSYPXR0) is invoked to determine if the message should be rerouted.

If the message is to be rerouted, the destination resolution exit (DFSYDRU0) is invoked to confirm the destination and prepare the header information. See Using OTMA exits in IMS and The pre-routing exit DFSYPXR0 for information about these exit programs.

Unless action is taken in the exits, all output from IMS transactions initiated from an IBM MQ queue manager, whether to the IOPCB or to an ALTPCB, will be returned to the same queue manager.

Sending unsolicited messages from IMS:

To send messages from IMS to an IBM MQ queue, you need to invoke an IMS transaction that ISRTs to an ALTPCB.

You need to write pre-routing and destination resolution exits to route unsolicited messages from IMS and build the OTMA user data, so that the MQMD of the message can be built correctly. See The pre-routing exit DFSYPXR0 and The destination resolution user exit for information about these exit programs.

Note: The IBM MQ - IMS bridge does not know whether a message that it receives is a reply or an unsolicited message. It handles the message the same way in each case, building the MQMD and MQIIH of the reply based on the OTMA UserData that arrived with the message.

Unsolicited messages can create new Tpipes. For example, if an existing IMS transaction switched to a new LTERM (for example PRINT01), but the implementation requires that the output be delivered through OTMA, a new Tpipe (called PRINT01 in this example) is created. By default, this is a non-synchronized Tpipe. If the implementation requires the message to be recoverable, set the destination resolution exit output flag. See the *IMS Customization Guide* for more information.

Message segmentation:

You can define IMS transactions as expecting single- or multi-segment input.

The originating IBM MQ application must construct the user input following the MQIIH structure as one or more LLZZ-data segments. All segments of an IMS message must be contained in a single IBM MQ message sent with a single MQPUT.

The maximum length of an LLZZ-data segment is defined by IMS/OTMA (32767 bytes). The total IBM MQ message length is the sum of the LL bytes, plus the length of the MQIIH structure.

All the segments of the reply are contained in a single IBM MQ message.

There is a further restriction on the 32 KB limitation on messages with format MQFMT_IMS_VAR_STRING. When the data in an ASCII-mixed CCSID message is converted to an EBCDIC-mixed CCSID message, a shift-in byte or a shift-out byte is added every time that there is a transition between SBCS and DBCS characters. The 32 KB restriction applies to the maximum size of the message. That is, because the LL field in the message cannot exceed 32 KB, the message must not exceed 32 KB including all shift-in and shift-out characters. The application building the message must allow for this.

Data conversion:

The data conversion is performed by either the distributed queuing facility (which may call any necessary exits) or by the intra group queuing agent (which does not support the use of exits) when it puts a message to a destination queue that has XCF information defined for its storage class.

Any exits needed must be available to the distributed queuing facility in the data set referenced by the CSQXLIB DD statement. This means that you can send messages to an IMS application using the IBM MQ - IMS bridge from any IBM MQ platform.

If there are conversion errors, the message is put to the queue unconverted; this results eventually in it being treated as an error by the IBM MQ - IMS bridge, because the bridge cannot recognize the header format. If a conversion error occurs, an error message is sent to the z/OS console.

See "Writing data-conversion exits" on page 364 for detailed information about data conversion in general.

Sending messages to the IBM MQ - IMS bridge

To ensure that conversion is performed correctly, you must tell the queue manager what the format of the message is.

If the message has an MQIIH structure, the *Format* in the MQMD must be set to the built-in format MQFMT_IMS, and the *Format* in the MQIIH must be set to the name of the format that describes your message data. If there is no MQIIH, set the *Format* in the MQMD to your format name.

If your data (other than the LLZZs) is all character data (MQCHAR), use as your format name (in the MQIIH or MQMD, as appropriate) the built-in format MQFMT_IMS_VAR_STRING. Otherwise, use your own format name, in which case you must also provide a data-conversion exit for your format. The exit must handle the conversion of the LLZZs in your message, in addition to the data itself (but it does not have to handle any MQIIH at the start of the message).

If your application uses *MFSMapName*, you can use messages with the MQFMT_IMS instead, and define the map name passed to the IMS transaction in the MFSMapName field of the MQIIH.

Receiving messages from the IBM MQ - IMS bridge

If an MQIIH structure is present on the original message that you are sending to IMS, one is also present on the reply message.

To ensure that your reply is converted correctly:

- If you have an MQIIH structure on your original message, specify the format that you want for your reply message in the MQIIH *ReplytoFormat* field of the original message. This value is placed in the MQIIH *Format* field of the reply message. This is particularly useful if all your output data is of the form LLZZ<character data>.
- If you do not have an MQIIH structure on your original message, specify the format that you want for the reply message as the MFS MOD name in the IMS application's ISRT to the IOPCB.

Developing MQI applications with IBM MQ

IBM MQ provides support for C, Visual Basic, COBOL, Assembler, RPG, pTAL, and PL/I. These procedural languages use the message queue interface (MQI) to access message queuing services.

For detailed information about how to write your applications in your chosen language, see the subtopics.

For an overview of the call interface for procedural languages, see Call descriptions. This topic contains a list of the MQI calls, and each call shows you how to code the calls in each of these languages.

IBM MQ provides data definition files to help you to write your applications. For a full description, see "IBM MQ data definition files" on page 71.

To help you choose which procedural language to code your programs in, consider the maximum length of the messages that your programs will process. If your programs will process only messages of a known maximum length, you can code them in any of the supported languages. If you do not know the maximum length of the messages that the programs will have to process, the language you choose will depend on whether you are writing a CICS, IMS, or batch application:

IMS and batch

Code the programs in C, PL/I, or assembler language to use the facilities these languages offer for obtaining and releasing arbitrary amounts of memory. Alternatively, you could code your programs in COBOL, but use assembler language, PL/I, or C subroutines to get and release storage.

CICS Code the programs in any language supported by CICS. The EXEC CICS interface provides the calls for managing memory, if necessary.

Related concepts:

“Object-oriented applications” on page 6

IBM MQ provides support for .NET, ActiveX, C++, Java, and JMS. These languages and frameworks use the IBM MQ Object Model, which provides classes that provide the same functionality as IBM MQ calls and structures. Some of the languages and frameworks that use the IBM MQ Object Model provide additional functions that are not available when you use procedural languages with the message queue interface (MQI).

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

Related information:

Technical overview

Application development reference

IBM MQ data definition files

IBM MQ provides data definition files to help you to write your applications.

Data definition files are also known as:

Language	Data definitions
C	Include files or header files
Visual Basic	Module files (32-bit versions only)
COBOL	Copy files
Assembler	Macros
PL/I	Include files

The data definition files to help you to write channel exits are described in IBM MQ COPY, header, include, and module files.

The data definition files to help you to write installable services exits are described in “User exits, API exits, and IBM MQ installable services” on page 311.

For data definition files supported on C++, see Using C++.



For data definition files supported on RPG, see the IBM i Application Programming Reference (ILE/RPG).



The names of the data definition files have the prefix CMQ, and a suffix that is determined by the programming language:

Suffix	Language
a	Assembler language
b	Visual Basic
c	C
l	COBOL (without initialized values)
p	PL/I
v	COBOL (with default values set)

Installation library



The name **thlqual** is the high-level qualifier of the installation library on z/OS.

This topic introduces IBM MQ data definition files, under these headings:

- “C language include files”
- “Visual Basic module files”
- “COBOL copy files”
-  “System/390 assembler-language macros” on page 74
-  “PL/I include files” on page 74

C language include files

The IBM MQ C include files are listed in C header files. They are installed in the following directories or libraries:

Platform	Installation directory or library
 IBM i	QMQM/H
UNIX	<i>MQ_INSTALLATION_PATH</i> /inc/
Windows systems	<i>MQ_INSTALLATION_PATH</i> \Tools\c\include
 z/OS	thlqual .SCSQC370

where *MQ_INSTALLATION_PATH* represents the high-level directory in which IBM MQ is installed.

Note: For UNIX, the include files are symbolically linked into */usr/include*.

For more information about the structure of directories, see Planning file system support.

Visual Basic module files

IBM MQ for Windows provides four Visual Basic module files.

They are listed in Visual Basic module files and installed in


MQ_INSTALLATION_PATH\Tools\Samples\VB\Include

COBOL copy files



For COBOL, IBM MQ provides separate copy files containing the named constants, and two copy files for each of the structures.

There are two copy files for each structure because each is provided both with and without initial values:

- In the WORKING-STORAGE SECTION of a COBOL program, use the files that initialize the structure fields to default values. These structures are defined in the copy files that have names suffixed with the letter V (values).
- In the LINKAGE SECTION of a COBOL program, use the structures without initial values. These structures are defined in copy files that have names suffixed with the letter L (linkage).

 Copy files containing data and interface definitions for IBM i are provided for ILE COBOL programs using prototyped calls to the MQI. The files exist in QMQM/QCBLLESRC with member names that have a suffix of L (for structures without initial values) or a suffix of V (for structures with initial values).

The IBM MQ COBOL copy files are listed in COBOL COPY files. They are installed in the following directories:

Platform	Installation directory or library
Other UNIX platforms	<i>MQ_INSTALLATION_PATH/inc/</i>
 IBM i	QMQM/QCBLLESRC
Windows	<i>MQ_INSTALLATION_PATH\Tools\cobol\copybook</i> (for Micro Focus COBOL) <i>MQ_INSTALLATION_PATH\Tools\cobol\copybook\VAcobol</i> (for IBM VisualAge® COBOL)
 z/OS	thlqual.SCSQCOBC

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

Include in your program only those files that you need. Do this with one or more COPY statements after a level-01 declaration. This means that you can include multiple versions of the structures in a program if necessary. Note that CMQV is a large file.

Here is an example of COBOL code to include the CMQMDV copy file:

```
01 MQM-MESSAGE-DESCRIPTOR.
COPY CMQMDV.
```

Each structure declaration begins with a level-01 item; you can declare several instances of the structure by coding the level-01 declaration followed by a COPY statement to copy in the remainder of the structure declaration. To refer to the appropriate instance, use the IN keyword.

Here is an example of COBOL code to include two instances of CMQMDV:

```
* Declare two instances of MQMD
01 MY-CMQMD.
COPY CMQMDV.
01 MY-OTHER-CMQMD.
COPY CMQMDV.
*
* Set MSGTYPE field in MY-OTHER-CMQMD
MOVE MQMT-REQUEST TO MQMD-MSGTYPE IN MY-OTHER-CMQMD.
```

Align the structures on 4-byte boundaries. If you use the COPY statement to include a structure following an item that is not the level-01 item, ensure that the structure is a multiple of 4-bytes from the start of the level-01 item. If you do not do this, you might reduce the performance of your application.

The structures are described in Data types used in the MQI. The descriptions of the fields in the structures show the names of fields without a prefix. In COBOL programs, prefix the field names with the name of the structure followed by a hyphen, as shown in the COBOL declarations. The fields in the structure copy files are prefixed in this way.

The field names in the declarations in the structure copy files are in uppercase. You can use mixed case or lowercase instead. For example, the field *StrucId* of the MQGMO structure is shown as MQGMO-STRUCID in the COBOL declaration and in the copy file.

The V-suffix structures are declared with initial values for all the fields, so you need to set only those fields where the value required is different from the initial value.

 z/OS

System/390 assembler-language macros

IBM MQ for z/OS provides two assembler-language macros containing the named constants, and one macro to generate each structure.

They are listed in z/OS Assembler COPY files and installed in **thlqual.SCSQMACS**.

These macros are called using code like this:

```
MY_MQMD CMQMDA EXPIRY=0,MSGTYPE=MQMT_DATAGRAM
```

▶ z/OS

PL/I include files

IBM MQ for z/OS provides include files that contain all the definitions that you need when you write IBM MQ applications in PL/I.

The files are listed in PL/I include files and installed in the **thlqual.SCSQPLIC** directory:

Include these files in your program if you are going to link the IBM MQ stub to your program (see “Preparing your program to run” on page 420). Include only CMQP if you intend to link the IBM MQ calls dynamically (see “Dynamically calling the IBM MQ stub” on page 427). Dynamic linking can be performed for batch and IMS programs only.

Writing a procedural application for queuing

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

Use the following links to find out more about writing applications:

- “The Message Queue Interface overview” on page 75
- “Connecting to and disconnecting from a queue manager” on page 91
- “Opening and closing objects” on page 100
- “Putting messages on a queue” on page 111
- “Getting messages from a queue” on page 128
- “Writing publish/subscribe applications” on page 172
- “Inquiring about and setting object attributes” on page 218
- “Committing and backing out units of work” on page 222
- “Starting IBM MQ applications using triggers” on page 234
- “Working with the MQI and clusters” on page 254
- ▶ z/OS “Using and writing applications on IBM MQ for z/OS” on page 258
- ▶ z/OS “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

Related concepts:

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

“Developing applications for IBM MQ” on page 1

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM MQ support applications written in procedural languages, and object oriented languages and frameworks.

“Design considerations for IBM MQ applications” on page 43

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

“Writing client procedural applications” on page 285

What you need to know to write client applications on IBM MQ using a procedural language.

“Building a procedural application” on page 382

You can write an IBM MQ application in one of several procedural languages, and run the application on several different platforms.

“Handling procedural program errors” on page 436

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Related tasks:

“Using the IBM MQ sample procedural programs” on page 457

These sample programs are written in procedural languages, and demonstrate typical uses of the Message Queue Interface (MQI). IBM MQ programs on different platforms.

“Developing web services with IBM MQ” on page 1217

You can develop IBM MQ applications for web services using the IBM MQ transport for SOAP.

The Message Queue Interface overview


Learn about the Message Queue Interface (MQI) components.

The Message Queue Interface consists of the following:

- *Calls* through which programs can access the queue manager and its facilities
- *Structures* that programs use to pass data to, and get data from, the queue manager
- *Elementary data types* for passing data to, and getting data from, the queue manager

 IBM MQ for z/OS also supplies:

- Two extra calls through which z/OS batch programs can commit and back out changes.
- *Data definition files* (sometimes known as copy files, macros, include files, and header files) that define the values of constants supplied with IBM MQ for z/OS.
- *Stub programs* to link-edit to your applications.
- A suite of sample programs that demonstrate how to use the MQI on the z/OS platform. For further information about these samples, see “Using the sample programs for z/OS” on page 568.


 IBM MQ for IBM i also supplies:

- *Data definition files* (sometimes known as copy files, macros, include files, and header files) that define the values of constants supplied with IBM MQ for IBM i.
- Three stub programs to link-edit to your ILE C, ILE COBOL, and ILE RPG applications.
- A suite of sample programs that demonstrate how to use the MQI on the IBM i platform.

IBM MQ for Windows and IBM MQ on UNIX and Linux systems also supply:

- Calls through which IBM MQ for Windows and IBM MQ on UNIX and Linux systems programs can commit and back out changes.
- *Include files* that define the values of constants supplied on these platforms.
- *Library files* to link your applications.
- A suite of sample programs that demonstrate how to use the MQI on these platforms. For further information about these samples, see “Using the sample programs on Multiplatforms” on page 457.
- Sample source and executable code for bindings to external transaction managers.

Use the following links to find out more about the MQI:

- “MQI calls”
- “Sync point calls” on page 78
- “Data conversion, data types, data definitions, and structures” on page 79
- “IBM MQ stub programs and library files” on page 80
- “Parameters common to all the calls” on page 86
- “Specifying buffers” on page 87
-  “z/OS batch considerations” on page 87
- “UNIX and Linux signal handling” on page 88

Related concepts:

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

 “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

 “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

This information helps you to write IMS applications using IBM MQ.

MQI calls:

Use this information to learn about calls in the Message Queue Interface (MQI).

The calls in the MQI can be grouped as follows:

MQCONN, MQCONNX, and MQDISC

Use these calls to connect a program to (with or without options), and disconnect a program

from, a queue manager. If you write CICS programs for z/OS, you do not need to use these calls. However, you are recommended to use them if you want to port your application to other platforms.

MQOPEN and MQCLOSE

Use these calls to open and close an object, such as a queue.

MQPUT and MQPUT1

Use these calls to put a message on a queue.

MQGET

Use this call to browse messages on a queue, or to remove messages from a queue.

MQSUB, MQSUBRQ

Use these calls to register a subscription to a topic, and to request publications matching the subscription.


MQINQ

Use this call to inquire about the attributes of an object.

MQSET

Use this call to set some of the attributes of a queue. You cannot set the attributes of other types of object.

MQBEGIN, MQCMIT, and MQBACK

Use these calls when IBM MQ is the coordinator of a unit of work. MQBEGIN starts the unit of work. MQCMIT and MQBACK end the unit of work, either committing or rolling back the updates made during the unit of work.  IBM i commitment controller is used to coordinate global units of work on IBM MQ for IBM i. Native start commitment control, commit, and rollback commands are used.

MQCRTMH, MQBUFMH, MQMHBUE, MQDLTMH

Use these calls to create a message handle, to convert a message handle to a buffer or a buffer to a message handle, and to delete a message handle.

MQSETMP, MQINQMP, MQDLTMP

Use these calls to set a message property on a message handle, inquire on a message property, and delete a property from a message handle.

MQCB, MQCB_FUNCTION, MQCTL

Use these calls to register and control a callback function.

MQSTAT

Use this call to retrieve status information about previous asynchronous put operations.

See Call descriptions for a description of the MQI calls.

Sync point calls:

Use this information to find out about sync point calls on different platforms.

Sync point calls are available as follows:

z/OS

IBM MQ for z/OS calls

IBM MQ for z/OS provides the MQCMIT and MQBACK calls.

Use these calls in z/OS batch programs to tell the queue manager that all the MQGET and MQPUT operations since the last sync point are to be made permanent (committed) or are to be backed out. To commit and back out changes in other environments:

CICS Use commands such as EXEC CICS SYNCPOINT and EXEC CICS SYNCPOINT ROLLBACK.

IMS Use the IMS sync point facilities, such as the GU (get unique) to the IOPCB, CHKP (checkpoint), and ROLB (rollback) calls.

RRS Use MQCMIT and MQBACK or SRRCMIT and SRRBACK as appropriate. (See “Transaction management and recoverable resource manager services” on page 226.)

Note: SRRCMIT and SRRBACK are native RRS commands, they are not MQI calls.

IBM i

IBM i calls

IBM MQ for IBM i provides the MQCMIT and MQBACK commands. You can also use the IBM i COMMIT and ROLLBACK commands, or any other commands or calls that initiate the IBM i commitment control facilities (for example, EXEC CICS SYNCPOINT).

ULW

IBM MQ calls on Windows, UNIX and Linux platforms

The following products provide the MQCMIT and MQBACK calls:

- IBM MQ for Windows
- IBM MQ on UNIX and Linux systems

Use sync point calls in programs to tell the queue manager that all the MQGET and MQPUT operations since the last sync point are to be made permanent (committed) or are to be backed out. To commit and back out changes in the CICS environment, use commands such as EXEC CICS SYNCPOINT and EXEC CICS SYNCPOINT ROLLBACK.

Data conversion, data types, data definitions, and structures:

Use this information to learn about data conversions, elementary data types, IBM MQ data definitions, and structures when using the Message Queue Interface.

Data conversion

The MQXCNVC (convert characters) call converts message character data from one character set to another. Except on IBM MQ for z/OS, this call is used only from a data-conversion exit.


See MQXCNVC - Convert characters for the syntax used with the MQXCNVC call, and “Writing data-conversion exits” on page 364 for guidance on writing and invoking data conversion exits.


Elementary data types

For the supported programming languages, the MQI provides elementary data types or unstructured fields.

These data types are described fully in Elementary data types.

IBM MQ data definitions

 IBM MQ for z/OS supplies data definitions in the form of COBOL copy files, assembly language macros, a single PL/I include file, a single C language include file, and C++ language include files.



 IBM MQ for IBM i supplies data definitions in the form of COBOL copy files, RPG copy files, C language include files, and C++ language include files.

The data definition files supplied with IBM MQ contain:

- Definitions of all the IBM MQ constants and return codes
- Definitions of the IBM MQ structures and data types
- Constant definitions for initializing the structures
- Function prototypes for each of the calls (for PL/I and the C language only)

For a full description of IBM MQ data definition files, see “IBM MQ data definition files” on page 71.

Structures

Structures, used with the MQI calls listed in “MQI calls” on page 76, are supplied in data definition files for each of the supported programming languages.   IBM MQ for z/OS and IBM MQ for IBM i supply files that contain constants for you to use when completing some of the fields of these structures. For more information about these, see IBM MQ data definitions.

See Structure data types summary for a summary of the structures.

IBM MQ stub programs and library files:

The stub programs and library files provided are listed here, for each platform.

For more information about how to use stub programs and library files when you build an executable application, see “Building a procedural application” on page 382. For information about linking to C++ library files, see *Using C++ IBM MQ Using C++*.

IBM MQ for AIX: 

On IBM MQ for AIX, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system.

In a non-threaded application, link to one of the following libraries:

Table 5. Library files for non-threaded AIX applications

Library file	Environment
libmqm.a	Server for C
libmqic.a & libmqm.a	Client for C
libmqmzf.a	Installable service exits for C
libmqmxa.a	Server XA interface
libmqmxa64.a	Server alternative XA interface
libmqcxa.a	Client XA interface
libmqcxa64.a	Client alternative XA interface
libmqmcbt.o	IBM MQ runtime library for Micro Focus COBOL support
libmqmcb.a	Server for COBOL
libmqicb.a	Client for COBOL
libimqc23ia.a	Client for C++
libimqs23ia.a	Server for C++

In a threaded application, link to one of the following libraries:

Table 6. Library files for threaded AIX applications.

A two column table listing the library files and the environment for each library file.

Library file	Environment
libmqm_r.a	Server for C
libmqic_r.a & libmqm_r.a	Client for C
libmqmzf_r.a	Installable service exits for C
libmqmxa_r.a	Server XA interface
libmqmxa64_r.a	Server alternative XA interface
libmqcxa_r.a	Client XA interface
libmqcxa64_r.a	Client alternative XA interface
libimqc23ia_r.a	Client for C++
libimqs23ia_r.a	Server for C++

Note: You cannot link to more than one library. That is, you cannot link to both a threaded and a non-threaded library at the same time.

IBM MQ for HP-UX: HP-UX

On IBM MQ for HP-UX, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system.

IA64 (IPF) platform

In a non-threaded application, link to one of the following libraries:

Table 7. Library files for non-threaded HP-UX applications

Library file	Environment
libmqm.so	Server for C
libmqic.so & libmqm.so	Client for C
libmqmzf.so	Installable service exits for C
libmqmxa.so	Server XA interface
libmqmxa64.so	Server alternative XA interface
libmqcxa.so	Client XA interface
libmqcxa64.so	Client alternative XA interface
libimqi23ah.so	C++
libmqmcbt.o	IBM MQ runtime library for Micro Focus COBOL support
libmqmcb.so	Server for COBOL
libmqicb.so	Client for COBOL

In a threaded application, link to one of the following libraries:

Table 8. Library files for threaded HP-UX applications

Library file	Environment
libmqm_r.so	Server for C
libmqmzf_r.so & libmqm_r.so	Installable service exits for C
libmqmxa_r.so	Server XA interface
libmqmxa64_r.so	Server alternative XA interface
libmqcxa_r.so	Client XA interface
libmqcxa64_r.so	Client alternative XA interface
libimqi23ah_r.so	C++

Note: You cannot link to more than one library. That is, you cannot link to both a threaded and a non-threaded library at the same time.

IBM MQ for IBM i: 

In IBM MQ for IBM i, link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system.

For non-threaded applications:

Table 9. Library files for non-threaded IBM i applications


Library file	Environment
LIBMQM	Server and Client service program
LIBMQIC	Client service program
IMQB23I4	C++ base service program
IMQS23I4	C++ server service program
LIBMQMZF	Installable exits for C

In a threaded application:

Table 10. Library files for threaded IBM i applications

Library file	Environment
LIBMQM_R	Server & client service program
IMQB23I4_R	C++ base service program
IMQS23I4_R	C++ server service program
LIBMQMZF_R	Installable exits for C
LIBMQIC_R	Client service program

On IBM MQ for IBM i, you can write your applications in C++. To see how to link your C++ applications, and for full details of all aspects of using C++, see Using C++.

IBM MQ for Linux: 

On IBM MQ for Linux, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system.

In a non-threaded application, link to one of the following libraries:

Table 11. Library files for non-threaded Linux applications

Library file	Environment
libmqm.so	Server for C
libmqic.so & libmqm.so	Client for C
libmqmzf.so	Installable service exits for C
libmqmxa.so	Server XA interface
libmqmxa64.so	Server alternative XA interface
libmqcxa.so	Client XA interface
libmqcxa64.so	Client alternative XA interface
libimqc23gl.so	Client for C++
libimqs23gl.so	Server for C++

In a threaded application, link to one of the following libraries:

Table 12. Library files for threaded Linux applications

Library file	Environment
libmqm_r.so	Server for C
libmqic_r.so & libmqm_r.so	Client for C
libmqmzf_r.so	Installable service exits for C
libmqmxa_r.so	Server XA interface
libmqmxa64_r.so	Server alternative XA interface
libmqcxa_r.so	Client XA interface
libmqcxa64_r.so	Client alternative XA interface
libimqc23gl_r.so	Client for C++
libimqs23gl_r.so	Server for C++

Note: You cannot link to more than one library. That is, you cannot link to both a threaded and a non-threaded library at the same time.

IBM MQ for Solaris: 

On IBM MQ for Solaris, you must link your program to the MQI library files supplied for the environment in which you are running your application in addition to those provided by the operating system.

Table 13. Library files for Solaris applications

Library File	Environment
libmqm.so	Server and client for C
libmqmzse.so	For C
libmqic.so	Client for C
libmqmcs.so	Common services for C
libmqmzf.so	Installable service exits for C
libmqmxa.so	Server XA interface
libmqmxa64.so	Server alternative XA interface
libmqcxa.so	Client XA interface
libmqcxa64.so	Client alternative XA interface
libimqc23as.a	Client for C++
libimqs23as.a	Server for C++

On IBM MQ for Windows, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system:

Table 14. Library files for Windows applications

Library File	Environment
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqm.lib</code>	Server for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqic.lib</code>	Client for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmxa.lib</code>	Server XA interface for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqcxa.lib</code>	Client XA interface for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqicxa.lib</code>	Client MTS for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcics4.lib32</code>	Server TXSeries CICS support for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqccics4.lib32</code>	Client TXSeries CICS support for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmzf.lib</code>	Installable services exits for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcbb.lib</code>	Server for IBM COBOL (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcb.lib</code>	Server for Micro Focus COBOL (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqicbb.lib</code>	Client for IBM COBOL (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqiccb.lib</code>	Client for Micro Focus COBOL (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\imqs23vn.lib</code>	Server for C++ (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\imqc23vn.lib</code>	Client for C++ (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\imqb23vn.lib</code>	Base for C++ (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\imqx23vn.lib</code>	Client MTS for C++ (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqm.lib</code>	Server for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqic.lib</code>	Client for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmxa.lib</code>	Server XA interface for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqcxa.lib</code>	Client XA interface for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqicxa.lib</code>	Client MTS for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmcbb.lib</code>	Server for IBM COBOL (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmcb.lib</code>	Server for Micro Focus COBOL (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqicbb.lib</code>	Client for IBM COBOL (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqiccb.lib</code>	Client for Micro Focus COBOL (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\imqs23vn.lib</code>	Server for C++ (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\imqc23vn.lib</code>	Client for C++ (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\imqb23vn.lib</code>	Base for C++ (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\imqx23vn.lib</code>	Client MTS for C++ (64-bit)


`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Use `amqmdnet.dll` for compiling .NET programs. See “Compiling IBM MQ .NET programs” on page 678 within the section “Developing .NET applications” on page 630 for more information.

These files are shipped for compatibility with previous releases:

`mqic32.lib`

mqic32xa.lib

IBM MQ for z/OS stub programs: 

Before you can run a program written with IBM MQ for z/OS, you must link-edit it to the stub program supplied with IBM MQ for z/OS for the environment in which you are running the application.

The stub program provides the first stage of the processing of your calls into requests that IBM MQ for z/OS can process.

IBM MQ for z/OS supplies the following stub programs:

CSQBSTUB

Stub program for z/OS batch programs

CSQBRRSI

Stub program for z/OS batch programs using RRS by way of the MQI

CSQBRSTB

Stub program for z/OS batch programs using RRS directly

CSQCSTUB

Stub program for CICS programs

CSQQSTUB

Stub program for IMS programs

CSQXSTUB

Stub program for distributed queuing non-CICS exits

CSQASTUB

Stub program for data-conversion exits

Attention: If you use a stub program other than one listed for a specific environment, it might have unpredictable results.

Note: If you use the CSQBRSTB stub program, link-edit with ATRSCSS from SYS1.CSSLIB. (SYS1.CSSLIB is also known as the *Callable Services Library*). For more information about RRS see “Transaction management and recoverable resource manager services” on page 226.

Alternatively, you can dynamically call the stub from within your program. This technique is described in “Dynamically calling the IBM MQ stub” on page 427.

In IMS, you might also need to use a special language interface module that is supplied by IBM MQ.

Do not run applications that are link-edited with CSQBSTUB and CSQQSTUB in the same IMS MPP region. This can cause problems such as DFS3607I or CSQQ005E messages. The first MQCONN call in an address space determines which interface is used, therefore CSQQSTUB and CSQBSTUB transactions must run in different IMS message regions.

Parameters common to all the calls:

There are two types of parameter common to all the calls: handles and return codes.

Using handles

All MQI calls use one or more *handles*. These identify the queue manager, queue or other object, message, or subscription, as appropriate to the call.

For a program to communicate with a queue manager, the program must have a unique identifier by which it knows that queue manager. This identifier is called a *connection handle*, sometimes referred to as a *Hconn*. For CICS programs, the connection handle is always zero. For all other platforms or styles of programs, the connection handle is returned by the MQCONN or MQCONNX call when the program connects to the queue manager. Programs pass the connection handle as an input parameter when they use the other calls.

For a program to work with an IBM MQ object, the program must have a unique identifier by which it knows that object. This identifier is called an *object handle*, sometimes referred to as an *Hobj*. The handle is returned by the MQOPEN call when the program opens the object to work with it. Programs pass the object handle as an input parameter when they use subsequent MQPUT, MQGET, MQINQ, MQSET, or MQCLOSE calls.

Similarly, the MQSUB call returns a *subscription handle* or *Hsub*, which is used to identify the subscription in subsequent MQGET, MQCB or MQSUBRQ calls, and certain calls processing message properties use a *message handle* or *Hmsg*.

Understanding return codes

A completion code and a reason code are returned as output parameters by each call. These are known collectively as *return codes*.

To show whether a call is successful, each call returns a *completion code* when the call is complete. The completion code is typically either MQCC_OK indicating success, or MQCC_FAILED indicating failure. Some calls can return an intermediate state, MQCC_WARNING, indicating partial success.

Each call also returns a *reason code* that shows the reason for the failure, or partial success, of the call. There are many reason codes, covering such circumstances as a queue being full, get operations not being allowed for a queue, and a particular queue not being defined for the queue manager. Programs can use the reason code to decide how to proceed. For example, they can prompt users to change their input data, then make the call again, or they can return an error message to the user.

When the completion code is MQCC_OK, the reason code is always MQRC_NONE.

The completion and reason codes for each call are listed with the description of that call. See Call descriptions and select the appropriate call from the list.

For more detailed information, including ideas for corrective action, see:

-  IBM MQ for z/OS messages, completion, and reason codes for IBM MQ for z/OS
- Reason codes for all other IBM MQ platforms

Specifying buffers:

The queue manager refers to buffers only if they are required. If you do not require a buffer on a call or the buffer is zero in length, you can use a null pointer to a buffer.

Always use datalength when specifying the size of the buffer that you require.

When you use a buffer to hold the output from a call (for example, to hold the message data for an MQGET call, or the values of attributes queried by the MQINQ call), the queue manager attempts to return a reason code if the buffer you specify is not valid or is in read-only storage. However, it might not always be able to return a reason code.

z/OS batch considerations:

z/OS batch programs that call the MQI can be in either supervisor or problem state.

However, they must meet the following conditions:

- They must be in task mode, not service request block (SRB) mode.
- They must be in Primary address space control (ASC) mode (not Access Register ASC mode).
- They must not be in cross-memory mode. The primary address space number (ASN) must be equal to the secondary ASN and the home ASN.
- They must not be used as MPF exit programs.
- No z/OS locks can be held.
- There can be no function recovery routines (FRRs) on the FRR stack.
- Any program status word (PSW) key can be in force for the MQCONN or MQCONNX call (provided the key is compatible with using storage that is in the TCB key), but subsequent calls that use the connection handle returned by MQCONN or MQCONNX:
 - Must have the same PSW key that was used on the MQCONN or MQCONNX call
 - Must have parameters accessible (for write, where appropriate) under the same PSW key
 - Must be issued under the same task (TCB), but not in any subtask of the task
- They can be in either 24-bit or 31-bit addressing mode. However, if 24-bit addressing mode is in force, parameter addresses must be interpreted as valid 31-bit addresses.

If any of these conditions is not met, a program check might occur. In some cases the call will fail and a reason code will be returned.

UNIX and Linux considerations:

Considerations that you need to be aware of.

Take note of the following points when developing UNIX and Linux applications.

The fork system call in UNIX and Linux systems:

UNIX

Linux

Note these considerations when using a fork system call in IBM MQ applications.

If your application wants to use fork, the parent process of that application should call fork before making any IBM MQ calls, for example, MQCONN, or creating an IBM MQ object using **ImqQueueManager**.

If your application wants to create a child process after making any IBM MQ calls, the application code must use a fork() with exec() to ensure that the child is a new instance, and not an exact copy of the parent.

If your application does not use exec(), the IBM MQ API call made within the child process returns MQRC_ENVIRONMENT_ERROR.

UNIX and Linux signal handling:

UNIX

Linux

This does not apply to IBM MQ for z/OS or IBM MQ for Windows.

In general, UNIX, Linux and IBM i systems have moved from a nonthreaded (process) environment to a multithreaded environment. In the nonthreaded environment, some functions could be implemented only by using signals, though most applications did not need to be aware of signals and signal handling. In the multithreaded environment, thread-based primitives support some of the functions that used to be implemented in the nonthreaded environments using signals.

In many instances, signals and signal handling, although supported, do not fit well into the multithreaded environment and various restrictions exist. This can be problematic when you are integrating application code with different middleware libraries (running as part of the application) in a multithreaded environment where each is trying to handle signals. The traditional approach of saving and restoring signal handlers (defined per process), which worked when there was only one thread of execution within a process, does not work in a multithreaded environment. This is because many threads of execution could be trying to save and restore a process-wide resource, with unpredictable results.

Unthreaded applications:

UNIX

Not applicable on Solaris as all applications are considered threaded even if they use only a single thread.

Each MQI function sets up its own signal handler for the signals:

- SIGALRM
- SIGBUS
- SIGFPE
- SIGSEGV
- SIGILL

Users' handlers for these are replaced for the duration of the MQI function call. Other signals can be caught in the normal way by user-written handlers. If you do not install a handler, the default actions (for example, ignore, core dump, or exit) are left in place.

After IBM MQ handles a synchronous signal (SIGSEGV, SIGBUS, SIGFPE, SIGILL), it attempts to pass the signal to any registered signal handler before making the MQI function call.

Threaded applications:  

A thread is considered to be connected to IBM MQ from MQCONN (or MQCONNX) until MQDISC.

Synchronous signals

Synchronous signals arise in a specific thread.

UNIX and Linux systems safely allow the setting up of a signal handler for such signals for the whole process. However, IBM MQ sets up its own handler for the following signals, in the application process, while any thread is connected to IBM MQ:

SIGBUS
SIGFPE
SIGSEGV
SIGILL

If you are writing multithreaded applications, there is only one process-wide signal handler for each signal. When IBM MQ sets up its own synchronous signal handlers it saves any previously registered handlers for each signal. After IBM MQ handles one of the signals listed, IBM MQ attempts to call the signal handler that was in effect at the time of the first IBM MQ connection within the process. The previously registered handlers are restored when all application threads have disconnected from IBM MQ.

Because signal handlers are saved and restored by IBM MQ, application threads must not establish signal handlers for these signals while there is any possibility that another thread of the same process is also connected to IBM MQ.

Note: When an application, or a middleware library (running as part of an application), establishes a signal handler while a thread is connected to IBM MQ, the application's signal handler must call the corresponding IBM MQ handler during the processing of that signal.

When establishing and restoring signal handlers, the general principle is that the last signal handler to be saved must be the first to be restored:

- When an application establishes a signal handler after connecting to IBM MQ, the previous signal handler must be restored before the application disconnects from IBM MQ.
- When an application establishes a signal handler before connecting to IBM MQ, the application must disconnect from IBM MQ before restoring its signal handler.

Note: Failure to observe the general principle that the last signal handler to be saved must be the first to be restored can result in unexpected signal handling in the application and, potentially, the loss of signals by the application.

Asynchronous signals

IBM MQ does not use any asynchronous signals in threaded applications unless they are client applications.

Additional considerations for threaded client applications

IBM MQ handles the following signals during I/O to a server. These signals are defined by the communications stack. The application must not establish a signal handler for these signals while a thread is connected to a queue manager:

SIGPIPE (for TCP/IP)

Additional considerations when using UNIX signal handling in MQI:

UNIX

Note these considerations when using UNIX signal handling.

Fastpath (trusted) applications

Fastpath applications run in the same process as IBM MQ and so are running in the multithreaded environment.

In this environment IBM MQ handles the synchronous signals SIGSEGV, SIGBUS, SIGFPE, and SIGILL. All other signals must not be delivered to the Fastpath application while it is connected to IBM MQ. Instead they must be blocked or handled by the application. If a Fastpath application intercepts such an event, the queue manager must be stopped and restarted, or it may be left in an undefined state. For a full list of the restrictions for Fastpath applications under MQCONN, see “Connecting to a queue manager using the MQCONN call” on page 94.

MQI function calls within signal handlers

While you are in a signal handler, do not call an MQI function.

If you try to call an MQI function from a signal handler while another MQI function is active, MQRC_CALL_IN_PROGRESS is returned. If you try to call an MQI function from a signal handler while no other MQI function is active, it is likely to fail sometime during the operation because of the operating system restrictions where only selective calls can be issued from, or within, a handler.

For C++ destructor methods, which might be called automatically during program exit, you might not be able to stop the MQI functions from being called. Ignore any errors about MQRC_CALL_IN_PROGRESS. If a signal handler calls exit(), IBM MQ backs out uncommitted messages in sync point as usual and closes any open queues.

Signals during MQI calls

MQI functions do not return the code EINTR or any equivalent to application programs.

If a signal occurs during an MQI call, and the handler calls *return*, the call continues to run as if the signal had not happened. In particular, MQGET cannot be interrupted by a signal to return control immediately to the application. If you want to break out of an MQGET, set the queue to GET_DISABLED; alternatively, use a loop around a call to MQGET with a finite time expiry (MQGMO_WAIT with gmo.WaitInterval set), and use your signal handler (in a nonthreaded environment) or equivalent function in a threaded environment to set a flag which breaks the loop.

AIX

In the AIX environment, IBM MQ requires that system calls interrupted by signals are restarted. When establishing your own signal handler with sigaction(2), set the SA_RESTART flag in the sa_flags field of the new action structure otherwise IBM MQ might be unable to complete any call interrupted by a signal.

User exits and installable services

User exits and installable services that run as part of an IBM MQ process in a multithreaded environment have the same restrictions as for fastpath applications. Consider these to be permanently connected to IBM MQ and so not using signals or non-threadsafe operating system calls.

VMS exit handlers

Users can install exit handlers for an IBM MQ application using the **SYSDCLEXH** system service.

The exit handler receives control when an image exits. An image exit normally occurs when you call the Exit (\$EXIT) or Force Exit (\$FORCEX) service. The \$FORCEX interrupts the target process in user mode. Then all user-mode exit handlers (established by \$DCLEXH) begin to execute in reverse order of establishment. For more details on exit handlers and \$FORCEX, see the *VMS Programming Concepts Manual* and the *VMS System Services Manual*.

If you call an MQI function from within an exit handler, the behavior of the function depends on the way the image was terminated. If the image was terminated while another MQI function is active, an MQRC_CALL_IN_PROGRESS is returned.

It is possible to call an MQI function from within an exit handler if no other MQI function is active and upcalls are disabled for the IBM MQ application. If upcalls are enabled for the IBM MQ application, it fails with the reason code MQRC_HCONN_ERROR.

The scope of an MQCONN or MQCONNX call is typically the thread that issued it. If upcalls are enabled, the exit handler runs as a separate thread and the connection handles cannot be shared.

Exit handlers are started within the interrupted context of the target process. It is up to the application to ensure that actions taken by a handler are safe and reliable, for the asynchronously interrupted context they are called from.

Connecting to and disconnecting from a queue manager

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

The way that this connection is made depends on the platform and the environment in which the program is operating:

Multi IBM MQ for Multiplatforms

Programs that run in these environments can use the MQCONN MQI call to connect to, and the MQDISC call to disconnect from, a queue manager. Alternatively, programs can use the MQCONNX call.

z/OS IBM MQ for z/OS batch

Programs that run in this environment can use the MQCONN MQI call to connect to, and the MQDISC call to disconnect from, a queue manager. Alternatively, programs can use the MQCONNX call.

z/OS batch programs can connect, consecutively or concurrently, to multiple queue managers on the same TCB.

z/OS IMS

The IMS control region is connected to one or more queue managers when it starts. This connection is controlled by IMS commands. For information about how to control the IMS adapter on z/OS, see *Administering IBM MQ for z/OS*. However, writers of message queuing IMS programs must use the MQCONN MQI call to specify the queue manager to which they want to connect. They can use the MQDISC call to disconnect from that queue manager.

Following an IMS call that establishes a syncpoint, and before processing a message for another user, the IMS adapter ensures that the application closes handles and disconnects from the queue manager. See “Syncpoints in IMS applications” on page 225.

IMS programs can connect, consecutively or concurrently, to multiple queue managers on the same TCB.

z/OS CICS Transaction Server for z/OS

CICS programs do not need to do any work to connect to a queue manager because the CICS system itself is connected. This connection is typically made automatically at initialization, but

you can also use the CKQC transaction that is supplied with IBM MQ for z/OS. For more information about CKQC, see *Administering IBM MQ for z/OS*.

CICS tasks can connect only to the queue manager that the CICS region is connected to.

CICS programs can also use the MQI connect and disconnect calls (MQCONN and MQDISC). You might want to do this so that you can port these applications to non-CICS environments with a minimum of recoding. However, these calls *always* complete successfully in a CICS environment. This means that the return code might not reflect the true state of the connection to the queue manager.

TXSeries for Windows and Open Systems

These programs do not need to do any work to connect to a queue manager because the CICS system itself is connected. Therefore, only one connection at a time is supported. CICS applications must issue an MQCONN call to obtain a connection handle, and an MQDISC call before they exit.

Use the following links to find out more about connecting and disconnecting from a queue manager:

- “Connecting to a queue manager using the MQCONN call”
- “Connecting to a queue manager using the MQCONN call” on page 94
- “Disconnecting programs from a queue manager using MQDISC” on page 99

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222


This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234


Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

 “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

 “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

This information helps you to write IMS applications using IBM MQ.

Connecting to a queue manager using the MQCONN call:

Use this information to learn how to connect to a queue manager using the MQCONN call.

In general, you can connect either to a specific queue manager, or to the default queue manager:

- For IBM MQ for z/OS, in the batch environment, the default queue manager is specified in the CSQBDEFV module.

- For IBM MQ for Windows, IBM i, UNIX, and Linux systems, the default queue manager is specified in the mq.ini file.

Alternatively, in the z/OS MVS batch, TSO, and RRS environments you can connect to any one queue manager within a queue-sharing group. The MQCONN or MQCONNX request selects any one of the active members of the group.

When you connect to a queue manager it must be local to the task. It must belong to the same system as the IBM MQ application.

In the IMS environment, the queue manager must be connected to the IMS control region and to the dependent region that the program uses. The default queue manager is specified in the CSQQDEFV module when IBM MQ for z/OS is installed.

With the TXSeries CICS environment, and TXSeries for Windows and AIX, the queue manager must be defined as an XA resource to CICS.

To connect to the default queue manager, call MQCONN, specifying a name consisting entirely of blanks or starting with a null (X'00') character.

An application must be authorized for it to successfully connect to a queue manager. For more information, see *Securing*.

The output from MQCONN is:

- A connection handle (**Hconn**)
- A completion code
- A reason code

Use the connection handle on subsequent MQI calls.

If the reason code indicates that the application is already connected to that queue manager, the connection handle that is returned is the same as the one that was returned when the application first connected. The application must not issue the MQDISC call in this situation because the calling application expects to remain connected.

The scope of the connection handle is the same as the scope of the object handle (see “Opening objects using the MQOPEN call” on page 101).

Descriptions of the parameters are given in the description of the MQCONN call in MQCONN.

The MQCONN call fails if the queue manager is in a quiescing state when you issue the call, or if the queue manager is shutting down.

Scope of MQCONN or MQCONNX

The scope of an MQCONN or MQCONNX call is typically the thread that issued it. That is, the connection handle returned from the call is valid only within the thread that issued the call. Only one call can be made at any one time using the handle. If it is used from a different thread, it is rejected as invalid. If you have multiple threads in your application and each wants to use IBM MQ calls, each one must issue MQCONN or MQCONNX.


It is not necessary for each call to be made to the same queue manager when a process makes multiple MQCONN calls. However, only one IBM MQ connection can be made from a thread at a time.

Alternatively, consider “Shared (thread independent) connections with MQCONNX” on page 97 to allow multiple IBM MQ connections from a single thread and an IBM MQ connection to be used from any thread.¹

If your application is running as a client, it can connect to more than one queue manager within a thread.

Connecting to a queue manager using the MQCONNX call:

The MQCONNX call is similar to the MQCONN call, but includes options to control the way that the call works.

As input to MQCONNX, you can supply a queue manager name  , or a queue-sharing group name on z/OS shared queue systems.

The output from MQCONNX is:

- A connection handle (Hconn)
- A completion code
- A reason code

You use the connection handle on subsequent MQI calls.

A description of all the parameters of MQCONNX is given in MQCONNX. The *Options* field allows you to set STANDARD_BINDING, FASTPATH_BINDING, SHARED_BINDING, or ISOLATED_BINDING for any version of MQCNO. You can also make shared (thread independent) connections using a MQCONNX call. See “Shared (thread independent) connections with MQCONNX” on page 97 for more information about these.

MQCNO_STANDARD_BINDING

By default, MQCONNX (like MQCONN) implies two logical threads where the IBM MQ application and the local queue manager agent run in separate processes. The IBM MQ application requests the IBM MQ operation and the local queue manager agent services the request. This is defined by the MQCNO_STANDARD_BINDING option on the MQCONNX call.

If you specify MQCNO_STANDARD_BINDING, the MQCONNX call uses either MQCNO_SHARED_BINDING or MQCNO_ISOLATED_BINDING, depending on the value of the DefaultBindType attribute of the queue manager, which is defined in qm.ini.

This is the default value.

If you are linking to the mqm library, then a standard server connection using the default bind type is attempted first. If the underlying server library failed to load, a client connection is attempted instead.

- If the MQ_CONNECT_TYPE environment variable is specified, one of the following options can be supplied to change the behavior of MQCONN or MQCONNX if MQCNO_STANDARD_BINDING is specified. (The exception to this is if MQCNO_FASTPATH_BINDING is specified with MQ_CONNECT_TYPE set to LOCAL or STANDARD to allow fastpath connections to be downgraded by the administrator without a related change to the application:

1. When using multithreaded applications with IBM MQ on UNIX and Linux systems you need to ensure that the applications have a sufficient stack size for the threads. Consider using a stack size of 256 KB, or larger, when multithreaded applications are making MQI calls, either by themselves or, with other signal handlers (for example, CICS).

Value	Meaning
CLIENT	A client connection only is attempted.
FASTPATH	This value was supported in previous releases, but is now ignored if specified.
LOCAL	A server connection only is attempted. Fastpath connections are downgraded to a standard server connection.
STANDARD	Supported for compatibility with previous releases. This value is now treated as LOCAL.

- If the MQ_CONNECT_TYPE environment variable is not set when MQCONN is called, a standard server connection using the default bind type is attempted. If the server library fails to load, a client connection is attempted.

MQCNO_FASTPATH_BINDING

Trusted applications imply that the IBM MQ application and the local queue manager agent become the same process. Because the agent process no longer needs to use an interface to access the queue manager, these applications become an extension of the queue manager. This is defined by the MQCNO_FASTPATH_BINDING option on the MQCONN call.

You need to link trusted applications to the threaded IBM MQ libraries. For instructions on how to set up an IBM MQ application to run as trusted, see MQCNO Options.

This option gives the highest performance.

Note: This option compromises the integrity of the queue manager: there is no protection from overwriting its storage. This also applies if the application contains errors that can be exposed to messages and other data in the queue manager too. Consider these issues before using this option.

MQCNO_SHARED_BINDING

Specify this option to make the application and the local queue manager agent run in separate processes. This maintains the integrity of the queue manager, that is, it protects the queue manager from errant programs. However, the application and the local queue manager agent share some resources.

This option is intermediate between MQCNO_FASTPATH_BINDING and MQCNO_ISOLATED_BINDING, both in terms of protecting the integrity of the queue manager, and in terms of the performance of MQI calls.

MQCNO_SHARED_BINDING is ignored if the queue manager does not support this type of binding. Processing continues as though the option had not been specified.

If an application has connected to the local queue manager using MQCNO_SHARED_BINDING, the queue manager can be stopped while the application is running. If you restart the queue manager while the application is still running, the attempt to start the queue manager fails with error AMQ7018 as the application is still holding on to resources needed by the queue manager.

In order to start the queue manager, you must stop the application.

MQCNO_ISOLATED_BINDING

Specify this option to make the application and the local queue manager agent run in separate processes, as for MQCNO_SHARED_BINDING. In this case, however, the application process and the local queue manager agent are isolated from each other in that they do not share resources.

This is the safest option for protecting the integrity of the queue manager, but it gives the slowest performance of MQI calls.

MQCNO_ISOLATED_BINDING is ignored if the queue manager does not support this type of binding. Processing continues as though the option had not been specified.

MQCNO_CLIENT_BINDING

Specify this option to make the application attempt a client connection only. This option has the following limitations:

- ▶ **z/OS** MQCNO_CLIENT_BINDING is ignored on z/OS.
- MQCNO_CLIENT_BINDING is rejected with MQRC_OPTIONS_ERROR if it is specified with any MQCNO binding option other than MQCNO_STANDARD_BINDING.
- MQCNO_CLIENT_BINDING is not available for Java as it has its own mechanisms for choosing the bind type.
- If the MQ_CONNECT_TYPE environment variable is not set when MQCONN is called, a standard server connection using the default bind type is attempted. If the server library fails to load, a client connection is attempted.

MQCNO_LOCAL_BINDING

Specify this option to make the application attempt a server connection. If either MQCNO_FASTPATH_BINDING, MQCNO_ISOLATED_BINDING, or MQCNO_SHARED_BINDING is also specified, then the connection is of that type instead, and is documented in this section. Otherwise a standard server connection is attempted using the default bind type. MQCNO_LOCAL_BINDING has the following limitations:

- ▶ **z/OS** MQCNO_LOCAL_BINDING is ignored on z/OS.
- MQCNO_LOCAL_BINDING is rejected with MQRC_OPTIONS_ERROR if it is specified with any MQCNO reconnect option other than MQCNO_RECONNECT_AS_DEF.
- MQCNO_LOCAL_BINDING is not available for Java as it has its own mechanisms for choosing the bind type.
- If the MQ_CONNECT_TYPE environment variable is not set when MQCONN is called, a standard server connection using the default bind type is attempted. If the server library fails to load, a client connection is attempted.

▶ **z/OS** On z/OS these options are tolerated, but only a standard bound connection is performed.

▶ **z/OS** MQCNO Version 3, for z/OS, allows four different options:

MQCNO_SERIALIZE_CONN_TAG_QSG

This allows an application to request that only one instance of an application runs at any one time in a queue-sharing group. This is achieved by registering the use of a connection tag with a value that is specified or derived by the application. The tag is a 128 byte character string specified in the Version 3 MQCNO.

MQCNO_RESTRICT_CONN_TAG_QSG

This is used where an application consists of more than one process (or a TCB), each of which can connect to a queue manager. Connection is permitted only if there is no current use of the tag, or the requesting application is within the same processing scope. This is MVS address space within the same queue-sharing group as the tag owner.

MQCNO_SERIALIZE_CONN_TAG_Q_MGR


This is similar to MQCNO_SERIALIZE_CONN_TAG_QSG, but only the local queue manager is interrogated to see if the requested tag is already in use.

MQCNO_RESTRICT_CONN_TAG_Q_MGR

This is similar to MQCNO_RESTRICT_CONN_TAG_QSG, but only the local queue manager is interrogated to see if the requested tag is already in use.

Restrictions for trusted applications:

The following restrictions apply to trusted applications:

- You must explicitly disconnect trusted applications from the queue manager.
- You must stop trusted applications before ending the queue manager with the `endmqm` command.
- You must not use asynchronous signals and timer interrupts (such as `sigkill`) with `MQCNO_FASTPATH_BINDING`.
- On all platforms, a thread within a trusted application cannot connect to a queue manager while another thread in the same process is connected to a different queue manager.
- On IBM MQ on UNIX and Linux systems you must use `mqm` as the effective `userID` and `groupID` for all MQI calls. You can change these IDs before making a non-MQI call requiring authentication (for example, opening a file), but you must change it back to `mqm` before making the next MQI call.
-  On IBM MQ for IBM i:
 1. Trusted applications must run under the QMQM user profile. It is not sufficient that the user profile be a member of the QMQM group or that the program adopt QMQM authority. It might not be possible for the QMQM user profile to be used to sign on to interactive jobs, or to be specified in the job description for jobs running trusted applications. In this case one approach is to use the IBM i profile swapping API functions, `QSYGETPH`, `QWTSETP`, and `QSYRLSPH` to temporarily change the current user of the job to QMQM while the MQ programs run. Details of these functions, together with an example of their use, is provided in the Security APIs section of the *IBM i System API Reference*.
 2. Do not cancel trusted applications using System-Request Option 2, or by ending the jobs in which they are running using `ENDJOB`.
- On IBM MQ for HP-UX, multithreaded fast-path applications are likely to need to set a larger stack size than the default. Use a size of 256 KB.
- On IBM MQ for Windows trusted 64-bit applications are not supported. If you try to run a trusted 64-bit application, it will be downgraded to a standard bound connection.
- On IBM MQ on UNIX and Linux systems trusted 32-bit applications are not supported. If you try to run a trusted 32-bit application, it will be downgraded to a standard bound connection.

Shared (thread independent) connections with MQCONN:

Use this information to learn about Shared connections with `MQCONN`, and some usage notes to consider.

Note: Not supported on IBM MQ for z/OS.

On IBM MQ platforms other than IBM MQ for z/OS, a connection made with `MQCONN` is available only to the thread that made the connection. Options on the `MQCONN` call allow you to create a connection that can be shared by all the threads in a process. If your application is running in a transactional environment that requires MQI calls to be issued on the same thread, you must use the following default option:

MQCNO_HANDLE_SHARE_NONE

Creates a non-shared connection.

In most other environments, you can use one of the following thread independent, shared connection options:

MQCNO_HANDLE_SHARE_BLOCK

Creates a shared connection. On a `MQCNO_HANDLE_SHARE_BLOCK` connection, if the connection is currently in use by an MQI call on another thread, the MQI call waits until the current MQI call has completed.

MQCNO_HANDLE_SHARE_NO_BLOCK

Creates a shared connection. On a MQCNO_HANDLE_SHARE_NO_BLOCK connection, if the connection is currently in use by an MQI call on another thread, the MQI call fails immediately with a reason of MQRC_CALL_IN_PROGRESS.

Except for the MTS (Microsoft Transaction Server) environment, the default value is MQCNO_HANDLE_SHARE_NONE. In the MTS environment, the default value is MQCNO_HANDLE_SHARE_BLOCK.

A connection handle is returned from the MQCONN call. The handle can be used by subsequent MQI calls from any thread in the process, associating those calls with the handle returned from the MQCONN. MQI calls using a single shared handle are serialized across threads.

For example, the following sequence of activity is possible with a shared handle:

1. Thread 1 issues MQCONN and gets a shared handle *h1*
2. Thread 1 opens a queue and issues a get request using *h1*
3. Thread 2 issues a put request using *h1*
4. Thread 3 issues a put request using *h1*
5. Thread 2 issues MQDISC using *h1*

While the handle is in use by any thread, access to the connection is unavailable to other threads. In circumstances where it is acceptable that a thread waits for any previous call from another thread to complete, use MQCONN with the option MQCNO_HANDLE_SHARE_BLOCK.

However blocking can cause difficulties. Suppose that in step 2, thread 1 issues a get request that waits for messages that might not have yet arrived (a get with wait). In this case, threads 2 and 3 are also left waiting (blocked) for as long as the get request on thread 1 takes. If you prefer that an MQI call returns with an error if another MQI call is already running on the handle, use MQCONN with the option MQCNO_HANDLE_SHARE_NO_BLOCK.

Shared connection usage notes

1. Any object handles (Hobj) created by opening an object are associated with an Hconn; so for a shared Hconn, the Hobjs are also shared and usable by any thread using the Hconn. Similarly, any unit of work started under an Hconn is associated with that Hconn; so this too is shared across threads with the shared Hconn.
2. *Any* thread can call MQDISC to disconnect a shared Hconn, not just the thread that called the corresponding MQCONN. The MQDISC terminates the Hconn making it unavailable to all threads.
3. A single thread can use multiple shared Hconns serially, for example use MQPUT to put one message under one shared Hconn then put another message using another shared Hconn, with each operation being under a different local unit of work.
4. Shared Hconns cannot be used within a global unit of work.

MQCONN environment variable:

Use this information to understand the different MQCONN call options and how they are used with MQ_CONNECT_TYPE. Note that MQ_CONNECT_TYPE only has any effect for STANDARD bindings. For other bindings, MQ_CONNECT_TYPE is ignored.

On IBM MQ for IBM i, IBM MQ for Windows, and IBM MQ on UNIX and Linux systems, you can use the environment variable, MQ_CONNECT_TYPE in combination with the type of binding specified in the *Options* field of the MQCNO structure used on an MQCONN call.

Table 15. The MQ_CONNECT_TYPE environment variable

MQCONN call option	MQ_CONNECT_TYPE environment variable	Result
STANDARD	UNDEFINED	STANDARD
STANDARD	STANDARD	STANDARD
STANDARD	FASTPATH	STANDARD
STANDARD	CLIENT	CLIENT
STANDARD	LOCAL	STANDARD

If MQCNO_STANDARD_BINDING is not specified, you can use MQCNO_NONE, which defaults to MQCNO_STANDARD_BINDING.

Disconnecting programs from a queue manager using MQDISC:

Use this information to learn about disconnecting programs from a queue manager using MQDISC.

When a program that has connected to a queue manager using the MQCONN or MQCONN call has finished all interaction with the queue manager, it breaks the connection using the MQDISC call, except:

- On CICS Transaction Server for z/OS applications, where the call is optional unless MQCONN was used and you want to drop the connection tag before the application ends.
- On IBM MQ for IBM i where, when you sign off from the operating system, an implicit MQDISC call is made.

As input to the MQDISC call, you must supply the connection handle (Hconn) that was returned by MQCONN or MQCONN when you connected to the queue manager.

Except on CICS on z/OS, after MQDISC is called the connection handle (Hconn) is no longer valid, and you cannot issue any further MQI calls until you call MQCONN or MQCONN again. MQDISC does an implicit MQCLOSE for any objects that are still open using this handle.

z/OS For a client connected to z/OS, when an MQDISC call is issued an implicit commit takes place, but any queue handles that are still open are not closed until the channel actually ends.

If you use MQCONN to connect on IBM MQ for z/OS, MQDISC also ends the scope of the connection tag established by the MQCONN. However, in a CICS, IMS, or RRS application, if there is an active unit of recovery associated with a connection tag, the MQDISC is rejected with a reason code of MQRC_CONN_TAG_NOT_RELEASED.

Descriptions of the parameters are given in the description of the MQDISC call in MQDISC.

When no MQDISC is issued

A standard, non-shared connection (Hconn) is cleaned up when the creating thread terminates. A shared connection is only implicitly backed out and disconnected when the whole process terminates. If the thread that created the shared Hconn terminates while the Hconn still exists the Hconn is still usable.

Authority checking

The MQCLOSE and MQDISC calls usually perform no authority checking.

In the normal course of events a job that has the authority to open or connect to an IBM MQ object closes or disconnect from that object. Even if the authority of a job that has connected to or opened an IBM MQ object is revoked, the MQCLOSE and MQDISC calls are accepted.

Opening and closing objects

This information provides an insight into opening and closing IBM MQ objects.

To perform any of the following operations, you must first *open* the relevant IBM MQ object:

- Put messages on a queue
- Get (browse or retrieve) messages from a queue
- Set the attributes of an object
- Inquire about the attributes of any object

Use the MQOPEN call to open the object, using the options of the call to specify what you want to do with the object. The only exception is if you want to put a single message on a queue, then close the queue immediately. In this case, you can bypass the *opening* stage by using the MQPUT1 call (see “Putting one message on a queue using the MQPUT1 call” on page 121).

Before you open an object using the MQOPEN call, you must connect your program to a queue manager. This is explained in detail, for all environments, in “Connecting to and disconnecting from a queue manager” on page 91.

There are four types of IBM MQ object that you can open:

- Queue
- Namelist
- Process definition
- Queue manager

You open all these objects in a similar way using the MQOPEN call. For more information about IBM MQ objects, see Object types.

You can open the same object more than once, and each time you get a new object handle. You might want to browse messages on a queue using one handle, and remove messages from the same queue using another handle. This saves using up resources to close and reopen the same object. You can also open a queue for browsing *and* removing messages at the same time.

Moreover, you can open multiple objects with a single MQOPEN and close them using MQCLOSE. See “Distribution lists” on page 122 for information about how to do this.

When you attempt to open an object, the queue manager checks that you are authorized to open that object for the options that you specify in the MQOPEN call.

Objects are closed automatically when a program disconnects from the queue manager. In the IMS environment, disconnection is forced when a program starts processing for a new user following a GU (get unique) IMS call. On the IBM i platform, objects are closed automatically when a job ends.

It is good programming practice to close objects you have opened. Use the MQCLOSE call to do this.

Use the following links to find out more about opening and closing objects:

- “Opening objects using the MQOPEN call”
- “Creating dynamic queues” on page 109
- “Opening remote queues” on page 110
- “Closing objects using the MQCLOSE call” on page 110

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

▶ **z/OS** “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

▶ **z/OS** “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

This information helps you to write IMS applications using IBM MQ.

Opening objects using the MQOPEN call:

Use this information to learn about opening objects using the MQOPEN call.

As input to the MQOPEN call, you must supply:

- A connection handle. For CICS applications on z/OS, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN or MQCONNX call. For other programs, always use the connection handle returned by the MQCONN or MQCONNX call.
- A description of the object that you want to open, using the object descriptor structure (MQOD).
- One or more options that control the action of the call.

The output from MQOPEN is:

- An object handle that represents your access to the object. Use this on input to any subsequent MQI calls.
- A modified object-descriptor structure, if you are creating a dynamic queue (and it is supported on your platform).
- A completion code.

- A reason code.

Scope of an object handle

The scope of an object handle (Hobj) is the same as the scope of a connection handle (Hconn).

This is covered in “Scope of MQCONN or MQCONNX” on page 93 and “Shared (thread independent) connections with MQCONNX” on page 97. However, there are additional considerations in some environments:

CICS In a CICS program, you can use the handle only within the same CICS task from which you made the MQOPEN call.

IMS and z/OS batch

In the IMS and batch environments, you can use the handle within the same task, but not within any subtasks.

Descriptions of the parameters of the MQOPEN call are given in MQOPEN.

The following sections describe the information that you must supply as input to MQOPEN.

Identifying objects (the MQOD structure)

Use the MQOD structure to identify the object that you want to open. This structure is an input parameter for the MQOPEN call. (The structure is modified by the queue manager when you use the MQOPEN call to create a dynamic queue.)

For full details of the MQOD structure, see MQOD.

For information about using the MQOD structure for distribution lists, see “Using the MQOD structure” on page 123 under “Distribution lists” on page 122.

Name resolution:

How the MQOPEN call resolves queue and queue manager names.

Note: A Queue manager alias is a remote queue definition without an RNAME field.

When you open an IBM MQ queue, the MQOPEN call performs a name resolution function on the queue name that you specify. This determines on which queue the queue manager performs subsequent operations. This means that when you specify the name of an alias queue or a remote queue in your object descriptor (MQOD), the call resolves the name either to a local queue or to a transmission queue. If a queue is opened for any type of input, browse, or set, it resolves to a local queue if there is one, and fails if there is not one. It resolves to a nonlocal queue only if it is opened for output only, inquire only, or output and inquire only. See Table 16 on page 103 for an overview of the name resolution process. The name that you supply in *ObjectQMGrName* is resolved *before* that in *ObjectName*.

Table 16 on page 103 also shows how you can use a local definition of a remote queue to define an alias for the name of a queue manager. This allows you to select which transmission queue is used when you put messages on a remote queue, so you could, for example, use a single transmission queue for messages destined for many remote queue managers.

To use the following table, first read down the two left-hand columns, under the heading **Input to MQOD**, and select the appropriate case. Then read across the corresponding row, following any instructions. Following the instructions in the **Resolved names** columns, you can either return to the **Input to MQOD** columns and insert values as directed, or you can exit the table with the results supplied. For example, you might be required to input *ObjectName*.

Table 16. Resolving queue names when using MQOPEN

Input to MQOD	Input to MQOD	Resolved names	Resolved names	Resolved names
<i>ObjectQMgrName</i>	<i>ObjectName</i>	<i>ObjectQMgrName</i>	<i>ObjectName</i>	Transmission queue
Blank or local queue manager	Local queue with no CLUSTER attribute	Local queue manager	Input <i>ObjectName</i>	Not applicable (local queue used)
Blank queue manager	Local queue with CLUSTER attribute	Workload management selected cluster queue manager or specific cluster queue manager selected on PUT	Input <i>ObjectName</i>	SYSTEM.CLUSTER.TRANSMIT.QUEUE and local queue used SYSTEM.QSG.TRANSMIT.QUEUE (see note)
Local queue manager	Local queue with CLUSTER attribute	Local queue manager	Input <i>ObjectName</i>	Not applicable (local queue used)
Blank or local queue manager	Model queue	Local queue manager	Generated name	Not applicable (local queue used)
Blank or local queue manager	Alias queue with or without CLUSTER attribute	Perform name resolution again with <i>ObjectQMgrName</i> unchanged, and input <i>ObjectName</i> set to the <i>BaseQName</i> in the alias queue definition object. Must not resolve to an alias locally defined where the <i>ObjectQMgrName</i> is specified, but can resolve to a clustered alias (hosted on other queue managers) where the <i>ObjectQMgrName</i> is blank.		
Local queue manager	Alias queue with CLUSTER attribute	The alias must not resolve to a cluster queue that is not locally defined, or a cluster queue that has the same <i>ObjectName</i> as the alias.		
Blank queue manager	Alias queue with CLUSTER attribute	The alias can resolve to a cluster queue with same <i>ObjectName</i> as the alias.		
Blank or local queue manager	Local definition of a remote queue	Perform name resolution again with <i>ObjectQMgrName</i> set to <i>RemoteQMgrName</i> , and <i>ObjectName</i> set to <i>RemoteQName</i> . Must not resolve remote queues		Name of <i>XmitQName</i> attribute, if non-blank; otherwise <i>RemoteQMgrName</i> in the remote queue definition object. SYSTEM.QSG.TRANSMIT.QUEUE (see note)

Table 16. Resolving queue names when using MQOPEN (continued)

Input to MQOD	Input to MQOD	Resolved names	Resolved names	Resolved names
Blank queue manager	No matching local object; cluster queue found	Workload management selected cluster queue manager or specific cluster queue manager selected on PUT	Input <i>ObjectName</i>	SYSTEM.CLUSTER. TRANSMIT.QUEUE SYSTEM.QSG. TRANSMIT.QUEUE (see note)
Blank or local queue manager	No matching local object; cluster queue not found		Error, queue not found	Not applicable
Name of queue manager in same queue-sharing group as local queue manager	Local shared queue	Local queue manager	Input <i>ObjectName</i>	Not applicable
Name of a local transmission queue	(Not resolved)	Input <i>ObjectQMgrName</i>	Input <i>ObjectName</i>	Input <i>ObjectQMgrName</i> SYSTEM.QSG. TRANSMIT.QUEUE (see note)
Queue manager alias definition (<i>RemoteQMgrName</i> may be the local queue manager)	(Not resolved, remote queue)	Perform name resolution again with <i>ObjectQMgrName</i> set to <i>RemoteQMgrName</i> . Must not resolve to remote queues	Input <i>ObjectName</i>	Name of <i>XmitQName</i> attribute, if non-blank; otherwise <i>RemoteQMgrName</i> in the remote queue definition object. SYSTEM.QSG. TRANSMIT.QUEUE (see note)
Queue manager is not the name of any local object; cluster queue managers or queue manager alias found	(Not resolved)	<i>ObjectQMgrName</i> or specific cluster queue manager selected on PUT	Input <i>ObjectName</i>	SYSTEM.CLUSTER. TRANSMIT.QUEUE SYSTEM.QSG. TRANSMIT.QUEUE (see note)
Queue manager is not the name of any local object; no cluster objects found	(Not resolved)	Input <i>ObjectQMgrName</i>	Input <i>ObjectName</i>	<i>DefXmitQName</i> attribute of the queue manager where <i>DefXmitQName</i> is supported. SYSTEM.QSG. TRANSMIT.QUEUE (see note)

Notes:

1. *BaseQName* is the name of the base queue from the definition of the alias queue.
2. *RemoteQName* is the name of the remote queue from the local definition of the remote queue.
3. *RemoteQMgrName* is the name of the remote queue manager from the local definition of the remote queue.
4. *XmitQName* is the name of the transmission queue from the local definition of the remote queue.
5. When using IBM MQ for z/OS queue managers that are part of a queue-sharing group (QSG), the name of the queue-sharing group can be used instead of the local queue manager name in Table 16 on page 103.

If the local queue manager cannot open the target queue, or put a message to the queue, the message is transferred to the specified `ObjectQMGrName` through, either intra-group queuing, or an IBM MQ channel.

6. In the *ObjectName* column of the table, CLUSTER refers to both the CLUSTER and CLUSNL attributes of the queue.
7. The `SYSTEM.QSG.TRANSMIT.QUEUE` is used if local and remote queue managers are in the same queue-sharing group; intra-group queuing is enabled.
8. If you have assigned a different cluster transmission queue to each cluster-sender channel, `SYSTEM.CLUSTER.TRANSMIT.QUEUE` might not be the name of the cluster transmission queue. For more information about multiple cluster transmission queues, see *Clustering: Planning how to configure cluster transmission queues*.
9. In the situation where the queue manager is not the name of any local object; cluster queue managers, or queue manager alias found.

When you have supplied a queue manager name using **ObjectQMGrName**, and there are multiple cluster channels with different cluster names known by the local queue manager that would reach that destination, then any of these channels might be used to move the message, regardless of the cluster name of the destination queue.

This might be unexpected, if you were anticipating messages for that queue only to be sent through a channel that has the same cluster name as the queue.

However, the **ObjectQMGrName** takes precedence in this case, and cluster workload balancing takes into consideration all channels that might reach that queue manager, regardless of the cluster name they are in.

Opening an alias queue also opens the base queue to which the alias resolves, and opening a remote queue also opens the transmission queue. Therefore you cannot delete either the queue that you specify or the queue to which it resolves while the other one is open.

While an alias queue is unable to resolve to another locally defined alias queue (shared in a cluster or not), resolving to a remotely defined cluster alias queue is permitted and can therefore be specified as the base queue.

The resolved queue name and the resolved queue manager name are stored in the *ResolvedQName* and *ResolvedQMGrName* fields in the MQOD.

For more information about name resolution in a distributed queuing environment see *What is queue name resolution?*.

Using the options of the MQOPEN call:

In the **Options** parameter of the MQOPEN call, you must choose one or more options to control the access that you are given to the object that you are opening. With these options you can:

- Open a queue and specify that all messages put to that queue must be directed to the same instance of it
- Open a queue to allow you to put messages on it
- Open a queue to allow you to browse messages on it
- Open a queue to allow you to remove messages from it
- Open an object to allow you to inquire about and set its attributes (but you can set the attributes of queues only)
- Open a topic or topic string to publish messages to it
- Associate context information with a message
- Nominate an alternative user identifier to be used for security checks
- Control the call if the queue manager is in a quiescing state

MQOPEN option for cluster queue:

The binding used for the queue handle is taken from the **DefBind** queue attribute, which can take the value `MQBND_BIND_ON_OPEN`, `MQBND_BIND_NOT_FIXED`, or `MQBND_BIND_ON_GROUP`.

To route all messages put to a queue using `MQPUT` to the same queue manager by the same route, use the `MQOO_BIND_ON_OPEN` option on the `MQOPEN` call.

To specify that a destination is to be selected at `MQPUT` time, that is, on a message-by-message basis, use the `MQOO_BIND_NOT_FIXED` option on the `MQOPEN` call.

To specify that all messages in a message groups put to a queue using `MQPUT` are allocated to the same destination instance, use 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.

If you do not specify any of these options the default, `MQOO_BIND_AS_Q_DEF`, is used.

If you specify the name of a queue manager in the `MQOD`, the queue at that queue manager is selected. If the queue manager name is blank, any instance can be selected. See “`MQOPEN` and clusters” on page 254 for more information.

If you open a cluster queue using a `QALIAS` definition, some queue attributes are defined by the alias queue, and not the base queue. Cluster attributes are among the attributes of the base queue definition that are overridden by the alias queue. For example, in the following snippet, the cluster queue is opened with `MQOO_BIND_NOT_FIXED` and not `MQOO_BIND_ON_OPEN`. The cluster queue definition is advertised throughout the cluster, the alias queue definition is local to the queue manager.

```
DEFINE QLOCAL(CLQ1) CLUSTER(MYCLUSTER) DEFBIND(OPEN) REPLACE
DEFINE QALIAS(ACLQ1) TARGQ(CLQ1) DEFBIND(NOTFIXED) REPLACE
```

MQOPEN option for putting messages:

To open a queue or topic to put messages on it, use the `MQOO_OUTPUT` option.

MQOPEN option for browsing messages:

To open a queue so that you can *browse* the messages on it, use the `MQOPEN` call with the `MQOO_BROWSE` option.

This creates a *browse cursor* that the queue manager uses to identify the next message on the queue. For more information, see “`Browsing messages on a queue`” on page 165.

Note:

1. You cannot browse messages on a remote queue; do not open a remote queue using the `MQOO_BROWSE` option.
2. You cannot specify this option when opening a distribution list. For further information about distribution lists, see “`Distribution lists`” on page 122.
3. Use the `MQOO_CO_OP` in conjunction with `MQOO_BROWSE` if you are using cooperative browsing; see `Options`

MQOPEN options for removing messages:

Three options control the opening of a queue to remove messages from it.

You can use only one of them in any MQOPEN call. These options define whether your program has exclusive or shared access to the queue. *Exclusive access* means that, until you close the queue, only you can remove messages from it. If another program attempts to open the queue to remove messages, its MQOPEN call fails. *Shared access* means that more than one program can remove messages from the queue.

The most advisable approach is to accept the type of access that was intended for the queue when the queue was defined. The queue definition involved the setting of the *Shareability* and the **DefInputOpenOption** attributes. To accept this access, use the MQOO_INPUT_AS_Q_DEF option. Refer to Table 17 to see how the setting of these attributes affects the type of access that you will be given when you use this option.

Table 17. How queue attributes and options of the MQOPEN call affect access to queues

Queue attributes		Type of access with MQOPEN options		
<i>Shareability</i>	<i>DefInputOpenOption</i>	AS_Q_DEF	SHARED	EXCLUSIVE
SHAREABLE	SHARED	shared	shared	exclusive
SHAREABLE	EXCLUSIVE	exclusive	shared	exclusive
NOT_SHAREABLE*	SHARED*	exclusive	exclusive	exclusive
NOT_SHAREABLE	EXCLUSIVE	exclusive	exclusive	exclusive

Note: * Although you can define a queue to have this combination of attributes, the default input open option is overridden by the shareability attribute.

Alternatively:

- If you know that your application can work successfully even if other programs can remove messages from the queue at the same time, use the MQOO_INPUT_SHARED option. Table 17 shows how, in some cases you will be given exclusive access to the queue, even with this option.
- If you know that your application can work successfully only if other programs are prevented from removing messages from the queue at the same time, use the MQOO_INPUT_EXCLUSIVE option.

Note:

1. You cannot remove messages from a remote queue. Therefore you cannot open a remote queue using any of the MQOO_INPUT_* options.
2. You cannot specify this option when opening a distribution list. For further information, see “Distribution lists” on page 122.

MQOPEN options for setting and inquiring about attributes:

To open a queue so that you can set its attributes, use the MQOO_SET option.

You cannot set the attributes of any other type of object (see "Inquiring about and setting object attributes" on page 218).

To open an object so that you can inquire about its attributes, use the MQOO_INQUIRE option.

Note: You cannot specify this option when opening a distribution list.

MQOPEN options relating to message context:

If you want to be able to associate context information with a message when you put it on a queue, you must use one of the message context options when you open the queue.

The options allow you to differentiate between context information that relates to the *user* who originated the message, and that which relates to the *application* that originated the message. Also, you can opt to set the context information when you put the message on the queue, or you can opt to have the context taken automatically from another queue handle.

Related concepts:

"Message context" on page 40

Message context information allows the application that retrieves the message to find out about the originator of the message.


"Controlling message context information" on page 118

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. You can use the options field in the MQPMO structure to control context information.

MQOPEN option for alternative user authority:

When you attempt to open an object using the MQOPEN call, the queue manager checks that you have the authority to open that object. If you are not authorized, the call fails.

However, server programs might want the queue manager to check the authorization of the user they are working for, rather than the server's own authorization. To do this, they must use the MQOO_ALTERNATE_USER_AUTHORITY option of the MQOPEN call, and specify the alternative user ID in the *AlternateUserId* field of the MQOD structure. Typically, the server would get the user ID from the context information in the message it is processing.

MQOPEN option for queue manager quiescing: 

If you use the MQOPEN call when the queue manager is in a quiescing state, the call might fail, depending on which environment you are using.

In the CICS environment on z/OS, if you use the MQOPEN call when the queue manager is in a quiescing state, the call always fails.

In other z/OS environments, IBM i, Windows systems and in UNIX and Linux systems environments, the call fails when the queue manager is quiescing only if you use the MQOO_FAIL_IF QUIESCING option of the MQOPEN call.

MQOPEN option for resolving local queue names:

When you open a local, alias or model queue, the local queue is returned.

However, when you open a remote queue or cluster queue, the *ResolvedQName* and *ResolvedQMGrName* fields of the MQOD structure are filled with the names of the remote queue and remote queue manager found in the remote queue definition, or with the chosen remote cluster queue.

Use the MQOO_RESOLVE_LOCAL_Q option of the MQOPEN call to fill the *ResolvedQName* in the MQOD structure with the name of the local queue that was opened. The *ResolvedQMGrName* is similarly filled with the name of the local queue manager hosting the local queue. This field is available only with Version 3 of the MQOD structure; if the structure is less than Version 3, MQOO_RESOLVE_LOCAL_Q is ignored without an error being returned.

If you specify MQOO_RESOLVE_LOCAL_Q when opening, for example, a remote queue, *ResolvedQName* is the name of the transmission queue to which messages will be put. *ResolvedQMGrName* is the name of the local queue manager hosting the transmission queue.

Creating dynamic queues:

Use a dynamic queue when you do not need the queue after your application ends.

For example, you could use a dynamic queue for your reply-to queue. You specify the name of the reply-to queue in the *ReplyToQ* field of the MQMD structure when you put a message on a queue (see “Defining messages using the MQMD structure” on page 113).

To create a dynamic queue, you use a template known as a model queue, together with the MQOPEN call. You create a model queue using the IBM MQ commands or the operations and control panels. The dynamic queue that you create takes the attributes of the model queue.

When you call MQOPEN, specify the name of the model queue in the *ObjectName* field of the MQOD structure. When the call completes, the *ObjectName* field is set to the name of the dynamic queue that is created. Also, the *ObjectQMGrName* field is set to the name of the local queue manager.

You can specify the name of the dynamic queue that you create in three ways:

- Give the full name that you want in the *DynamicQName* field of the MQOD structure.
- Specify a prefix (fewer than 33 characters) for the name, and allow the queue manager to generate the rest of the name. This means that the queue manager generates a unique name, but you still have some control (for example, you might want each user to use a certain prefix, or you might want to give a special security classification to queues with a certain prefix in their name). To use this method, specify an asterisk (*) for the last non-blank character of the *DynamicQName* field. Do not specify a single asterisk (*) for the dynamic queue name.

- Allow the queue manager to generate the full name. To use this method, specify an asterisk (*) in the first character position of the *DynamicQName* field.

For more information about these methods, see the description of the *DynamicQName* field.

There is more information on dynamic queues in Dynamic and Model queues.

Opening remote queues:

A remote queue is a queue that is owned by a queue manager other than the one to which the application is connected.

To open a remote queue, use the MQOPEN call as for a local queue. You can specify the name of the queue as follows:

1. In the *ObjectName* field of the MQOD structure, specify the name of the remote queue as known to the *local* queue manager.

Note: Leave the *ObjectQMgrName* field blank in this case.

2. In the *ObjectName* field of the MQOD structure, specify the name of the remote queue, as known to the *remote* queue manager. In the *ObjectQMgrName* field, specify either:
 - The name of the transmission queue that has the same name as the remote queue manager. The name and case (uppercase, lowercase or a mixture) must match *exactly*.
 - The name of a queue manager alias object that resolves to the destination queue manager or the transmission queue.

This tells the queue manager the destination of the message as well as the transmission queue that it needs to be put on to get there.

3. If *DefXmitQname* is supported, in the *ObjectName* field of the MQOD structure, specify the name of the remote queue as known by the *remote* queue manager.

Note: Set the *ObjectQMgrName* field to the name of the remote queue manager (it cannot be left blank in this case).

Only local names are validated when you call MQOPEN; the last check is for the existence of the transmission queue to be used.

These methods are summarized in Table 16 on page 103.

Closing objects using the MQCLOSE call:

To close an object, use the MQCLOSE call.

If the object is a queue, note the following:

- You do not need to empty a temporary dynamic queue before you close it.
When you close a temporary dynamic queue, the queue is deleted, along with any messages that might still be on it. This is true even if there are uncommitted MQGET, MQPUT, or MQPUT1 calls outstanding against the queue.
- On IBM MQ for z/OS, if you have any MQGET requests with an MQGMO_SET_SIGNAL option outstanding for that queue, they are canceled.
- If you opened the queue using the MQOO_BROWSE option, your browse cursor is destroyed.

Closure is unrelated to sync point, so you can close queues before or after sync point.

As input to the MQCLOSE call, you must supply:

- A connection handle. Use the same connection handle used to open it, or alternatively, for CICS applications on z/OS, you can specify the constant MQHC_DEF_HCONN (which has the value zero).
- The handle of the object that you want to close. Get this from the output of the MQOPEN call.
- MQCO_NONE in the *options* field (unless you are closing a permanent dynamic queue).
- The control option to determine whether the queue manager should delete the queue even if there are still messages on it (when closing a permanent dynamic queue).

The output from MQCLOSE is:

- A completion code
- A reason code
- The object handle, reset to the value MQHO_UNUSABLE_HOBJ

Descriptions of the parameters of the MQCLOSE call are given in MQCLOSE.

Putting messages on a queue

Use this information to learn how to put messages on a queue.

Use the MQPUT call to put messages on the queue. You can use MQPUT repeatedly to put many messages on the same queue, following the initial MQOPEN call. Call MQCLOSE when you have finished putting all your messages on the queue.

If you want to put a single message on a queue and close the queue immediately afterward, you can use the MQPUT1 call. MQPUT1 performs the same functions as the following sequence of calls:

- MQOPEN
- MQPUT
- MQCLOSE

Generally however, if you have more than one message to put on the queue, it is more efficient to use the MQPUT call. This depends on the size of the message and the platform that you are working on.

Use the following links to find out more about putting messages on a queue:

- “Putting messages on a local queue using the MQPUT call” on page 112
- “Putting messages on a remote queue” on page 118
- “Setting properties of a message” on page 118
- “Controlling message context information” on page 118
- “Putting one message on a queue using the MQPUT1 call” on page 121
- “Distribution lists” on page 122
- “Some cases where the put calls fail” on page 127

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

▶ **z/OS** “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

▶ **z/OS** “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

This information helps you to write IMS applications using IBM MQ.

Putting messages on a local queue using the MQPUT call:

Use this information to learn about putting messages on a local queue using the MQPUT call.

As input to the MQPUT call, you must supply:

- A connection handle (Hconn).
- A queue handle (Hobj).
- A description of the message that you want to put on the queue. This is in the form of a message descriptor structure (MQMD).
- Control information, in the form of a put-message options structure (MQPMO).
- The length of the data contained within the message (MQLONG).
- The message data itself.

The output from the MQPUT call is as follows:

- A reason code (MQLONG)
- A completion code (MQLONG)

If the call completes successfully, it also returns your options structure and your message descriptor structure. The call modifies your options structure to show the name of the queue and the queue manager to which the message was sent. If you request that the queue manager generates a unique value for the identifier of the message you are putting (by specifying binary zero in the *MsgId* field of the MQMD structure), the call inserts the value in the *MsgId* field before returning this structure to you. Reset this value before you issue another MQPUT.

There is a description of the MQPUT call in MQPUT.

For more description on the information needed as input to the MQPUT call, see the following links:

- “Specifying handles”
- “Defining messages using the MQMD structure”
- “Specifying options using the MQPMO structure”
- “The data in your message” on page 116
- “Putting messages: Using message handles” on page 117

Specifying handles

For the connection handle (*Hconn*) in CICS on z/OS applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or you can use the connection handle returned by the MQCONN or MQCONNX call. For other applications, always use the connection handle returned by the MQCONN or MQCONNX call.

Whatever environment you are working in, use the same queue handle (*Hobj*) that is returned by the MQOPEN call.

Defining messages using the MQMD structure

The message descriptor structure (MQMD) is an input/output parameter for the MQPUT and MQPUT1 calls. Use it to define the message you are putting on a queue.

If MQPRI_PRIORITY_AS_Q_DEF or MQPER_PERSISTENCE_AS_Q_DEF is specified for the message and the queue is a cluster queue, the values used are those of the queue to which the MQPUT resolves. If that queue is disabled for MQPUT, the call will fail. See *Configuring a queue manager cluster* for more information.

Note: Use MQPMO_NEW_MSG_ID and MQPMO_NEW_CORREL_ID before putting a new message to ensure that the *MsgId* and *CorrelId* are unique. The values in these fields are returned on a successful MQPUT.

There is an introduction to the message properties that MQMD describes in “IBM MQ messages” on page 9, and there is a description of the structure itself in MQMD.

Specifying options using the MQPMO structure

Use the MQPMO (Put Message Option) structure to pass options to the MQPUT and MQPUT1 calls.

The following sections give you help on filling in the fields of this structure. There is a description of the structure in MQPMO.

The structure includes the following fields:

- *StrucId*
- *Version*
- *Options*
- *Context*
- *ResolvedQName*
- *ResolvedQMGrName*
- *RecsPresent*
- *PutMsgRecsFields*
- *ResponseRecOffset* and *ResponseRecPtr*
- *OriginalMsgHandle*

- *NewMsgHandle*
- *Action*
- *PubLevel*

The contents of these fields is as follows:

StrucId

This identifies the structure as a put-message options structure. This is a 4-character field. Always specify MQPMO_STRUC_ID.

Version

This describes the version number of the structure. The default is MQPMO_VERSION_1. If you enter MQPMO_VERSION_2, you can use distribution lists (see “Distribution lists” on page 122). If you enter MQPMO_VERSION_3, you can use message handles and message properties. If you enter MQPMO_CURRENT_VERSION, your application is set always to use the most recent level.

Options

This controls the following:

- Whether the put operation is included in a unit of work
- How much context information is associated with a message
- Where the context information is taken from
- Whether the call fails if the queue manager is in a quiescing state
- Whether grouping or segmentation is allowed
- Generation of a new message identifier and correlation identifier
- The order in which messages and segments are put on a queue
- Whether to resolve local queue names

If you leave the *Options* field set to the default value (MQPMO_NONE), the message you put has default context information associated with it.

Also, the way that the call operates with sync points is determined by the platform. The sync point control default is yes in z/OS ; for other platforms, it is no.

Context

This states the name of the queue handle that you want context information to be copied from (if requested in the *Options* field).

For an introduction to message context, see “Message context” on page 40. For information about using the MQPMO structure to control the context information in a message, see “Controlling message context information” on page 118.

ResolvedQName

This contains the name (after resolution of any alias name) of the queue that was opened to receive the message. This is an output field.

ResolvedQMgrName

This contains the name (after resolution of any alias name) of the queue manager that owns the queue in *ResolvedQName*. This is an output field.

The MQPMO can also accommodate fields required for distribution lists (see “Distribution lists” on page 122). If you want to use this facility, Version 2 of the MQPMO structure is used. This includes the following fields:

RecsPresent

This field contains the number of queues in the distribution list; that is, the number of Put Message Records (MQPMR) and corresponding Response Records (MQRR) present.

The value that you enter can be the same as the number of Object Records provided at MQOPEN. However, if the value is less than the number of Object Records provided on the

MQOPEN call, or if you provide no Put Message Records, the values of the queues that are not defined are taken from the default values provided by the message descriptor. Also, if the value is greater than the number of Object Records provided, the excess Put Message Records are ignored.

You are recommended to do one of the following:

- If you want to receive a report or reply from each destination, enter the same value as appears in the MQOR structure and use MQPMRs containing *MsgId* fields. Either initialize these *MsgId* fields to zeros or specify MQPMO_NEW_MSG_ID.

When you have put the message to the queue, *MsgId* values that the queue manager has created become available in the MQPMRs; you can use these to identify which destination is associated with each report or reply.

- If you do not want to receive reports or replies, choose one of the following:
 1. If you want to identify destinations that fail immediately, you might still want to enter the same value in the *RecsPresent* field as appears in the MQOR structure and provide MQRRs to identify these destinations. Do not specify any MQPMRs.
 2. If you do not want to identify failed destinations, enter zero in the *RecsPresent* field and do not provide MQPMRs nor MQRRs.

Note: If you are using MQPUT1, the number of Response Record Pointers and Response Record Offsets must be zero.

For a full description of Put Message Records (MQPMR) and Response Records (MQRR), see MQPMR and MQRR.

PutMsgRecFields

This indicates which fields are present in each Put Message Record (MQPMR). For a list of these fields, see “Using the MQPMR structure” on page 126.

PutMsgRecOffset and *PutMsgRecPtr*

Pointers (typically in C) and offsets (typically in COBOL) are used to address the Put Message Records (see “Using the MQPMR structure” on page 126 for an overview of the MQPMR structure).

Use the *PutMsgRecPtr* field to specify a pointer to the first Put Message Record, or the *PutMsgRecOffset* field to specify the offset of the first Put Message Record. This is the offset from the start of the MQPMO. Depending on the *PutMsgRecFields* field, enter a nonnull value for either *PutMsgRecOffset* or *PutMsgRecPtr*.

ResponseRecOffset and ResponseRecPtr

You also use pointers and offsets to address the Response Records (see “Using the MQRR structure” on page 125 for further information about Response Records).

Use the *ResponseRecPtr* field to specify a pointer to the first Response Record, or the *ResponseRecOffset* field to specify the offset of the first Response Record. This is the offset from the start of the MQPMO structure. Enter a nonnull value for either *ResponseRecOffset* or *ResponseRecPtr*.

Note: If you are using MQPUT1 to put messages to a distribution list, *ResponseRecPtr* must be null or zero and *ResponseRecOffset* must be zero.

Version 3 of the MQPMO structure additionally includes the following fields:

OriginalMsgHandle

The use you can make of this field depends on the value of the *Action* field. If you are putting a new message with associated message properties, set this field to the message handle you previously created and set properties on. If you are forwarding, replying to, or generating a report in response to a previously retrieved message, this field contains the message handle of that message.

NewMsgHandle

If you specify a *NewMsgHandle*, any properties associated with the handle override properties associated with the *OriginalMsgHandle*. For more information, see Action (MQLONG).

Action

Use this field to specify the type of put being performed. Possible values and their meanings are as follows:

MQACTP_NEW

This is a new message unrelated to any other.

MQACTP_FORWARD

This message was retrieved previously and is now being forwarded.

MQACTP_REPLY

This message is a reply to a previously retrieved message.

MQACTP_REPORT

This message is a report generated as a result of a previously retrieved message.

For more information, see Action (MQLONG).

PubLevel

If this message is a publication, you can set this field to determine which subscriptions receive it. Only subscriptions with a *SubLevel* less than or equal to this value will receive this publication. The default value is 9 which is the highest level and means that subscriptions with any *SubLevel* can receive this publication.

The data in your message

Give the address of the buffer that contains your data in the **Buffer** parameter of the MQPUT call. You can include anything in the data in your messages. The amount of data in the messages, however, affects the performance of the application that is processing them.

The maximum size of the data is determined by:

- The **MaxMsgLength** attribute of the queue manager
- The **MaxMsgLength** attribute of the queue on which you are putting the message
- The size of any message header added by IBM MQ (including the dead-letter header, MQDLH and the distribution list header, MQDH)

The **MaxMsgLength** attribute of the queue manager holds the size of message that the queue manager can process. This has a default of 100 MB for all IBM MQ products at V6 or higher.

To determine the value of this attribute, use the MQINQ call on the queue manager object. For large messages, you can change this value.

The **MaxMsgLength** attribute of a queue determines the maximum size of message that you can put on the queue. If you attempt to put a message with a size larger than the value of this attribute, your MQPUT call fails. If you are putting a message on a remote queue, the maximum size of message that you can successfully put is determined by the **MaxMsgLength** attribute of the remote queue, of any intermediate transmission queues that the message is put on along the route to its destination, and of the channels used.

For an MQPUT operation, the size of the message must be smaller than or equal to the **MaxMsgLength** attribute of both the queue and the queue manager. The values of these attributes are independent, but you are recommended to set the *MaxMsgLength* of the queue to a value less than or equal to that of the queue manager.

IBM MQ adds header information to messages in the following circumstances:


- When you put a message on a remote queue, IBM MQ adds a transmission header structure (MQXQH) to the message. This structure includes the name of the destination queue and its owning queue manager.
- If IBM MQ cannot deliver a message to a remote queue, it attempts to put the message on the dead-letter (undelivered-message) queue. It adds an MQDLH structure to the message. This structure includes the name of the destination queue and the reason that the message was put on the dead-letter queue.
- If you want to send a message to multiple destination queues, IBM MQ adds an MQDH header to the message. This describes the data that is present in a message, belonging to a distribution list, on a transmission queue. Consider this when choosing an optimum value for the maximum message length.
- If the message is a segment or a message in a group, IBM MQ might add an MQMDE.

These structures are described in MQDH and MQMDE.

If your messages are of the maximum size allowed for these queues, the addition of these headers means that the put operations fail because the messages are now too big. To reduce the possibility of the put operations failing:

- Make the size of your messages smaller than the **MaxMsgLength** attribute of the transmission and dead-letter queues. Allow at least the value of the MQ_MSG_HEADER_LENGTH constant (more for large distribution lists).
- Make sure that the **MaxMsgLength** attribute of the dead-letter queue is set to the same as the *MaxMsgLength* of the queue manager that owns the dead-letter queue.

The attributes for the queue manager and the message queuing constants are described in Attributes for the queue manager.

 For information on how undelivered messages are handled in a distributed queuing environment, see Undelivered/unprocessed messages.

Putting messages: Using message handles

Two message handles are available in the MQPMO structure, *OriginalMsgHandle* and *NewMsgHandle*. The relationship between these message handles is defined by the value of the MQPMO *Action* field.

For full details see Action (MQLONG). A message handle is not necessarily required in order to put a message. Its purpose is to associate properties with a message, so it is required only if you are using message properties.

Putting messages on a remote queue:

When you want to put a message on a remote queue (that is, a queue owned by a queue manager other than the one to which your application is connected) rather than a local queue, the only extra consideration is how you specify the name of the queue when you open it.

This is described in “Opening remote queues” on page 110. There is no change to how you use the MQPUT or MQPUT1 call for a local queue.

For more information on using remote and transmission queues, see IBM MQ distributed queuing techniques.

Setting properties of a message:

Call MQSETMP for each property you want to set. When you put the message set the message handle and action fields of the MQPMO structure.

To associate properties with a message, the message must have a message handle. Create a message handle using the MQCRTMH function call. Call MQSETMP specifying this message handle for each property you want to set. A sample program, amqsstma.c, is provided to illustrate the use of MQSETMP.

If this is a new message, when you put it to a queue, using MQPUT or MQPUT1, set the OriginalMsgHandle field in the MQPMO to the value of this message handle, and set the MQPMO Action field to MQACTP_NEW (this is the default value).

If this is a message you have previously retrieved, and you are now forwarding or replying to it or sending a report in response to it, put the original message handle in the OriginalMsgHandle field of the MQPMO and the new message handle in the NewMsgHandle field. Set the Action field to MQACTP_FORWARD, MQACTP_REPLY, or MQACTP_REPORT, as appropriate.

If you have properties in an MQRFH2 header from a message you have previously retrieved, you can convert them to message handle properties using the MQBUFMH call.

If you are putting your message to a queue on a queue manager at a level earlier than IBM WebSphere MQ Version 7.0, which cannot process message properties, you can set the PropertyControl parameter in the channel definition to specify how the properties are to be treated.

Controlling message context information:

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. You can use the options field in the MQPMO structure to control context information.

Message context information allows the application that retrieves the message to find out about the originator of the message. All context information is stored in the context fields of the message descriptor. The type of information falls into identity, origin, and user context information.

To control context information, use the *Options* field in the MQPMO structure.

If you do not specify any options for context information, the queue manager overwrites context information that might already be in the message descriptor with the identity and context information that it has generated for your message. This is the same as specifying the MQPMO_DEFAULT_CONTEXT option. You might want this default context information when you create a new message (for example, when processing user input from an inquiry screen).

If you want no context information associated with your message, use the `MQPMO_NO_CONTEXT` option. When putting a message with no context, any authority checks made by IBM MQ are made using a blank user ID. A blank user ID cannot be assigned explicit authority to IBM MQ resources but is treated as a member of the special group 'nobody'. For more details on the special group nobody, see Installable services interface reference information.

You can do context setting using `MQOPEN` followed by `MQPUT` using the `MQOO_` option and `MQPMO_` option indicated in the following sections. You can also do context setting using just an `MQPUT1`, in which case you just need to select the `MQPMO_` option indicated in the sections below.

The following sections of this topic explain the use of identity context, user context, and all context.

- “Passing identity context”
- “Passing user context”
- “Passing all context” on page 120
- “Setting identity context” on page 120
- “Setting user context” on page 120
- “Setting all context” on page 120

Passing identity context

In general, programs should pass identity context information from message to message around an application until the data reaches its final destination.

Programs should change the origin context information each time that they change the data. However, applications that want to change or set any context information must have the appropriate level of authority. The queue manager checks this authority when the applications open the queues; they must have authority to use the appropriate context options for the `MQOPEN` call.

If your application gets a message, processes the data from the message, then puts the changed data into another message (possibly for processing by another application), the application must pass the identity context information from the original message to the new message. You can allow the queue manager to create the origin context information.

To save the context information from the original message, use the `MQOO_SAVE_ALL_CONTEXT` option when you open the queue for getting the message. This is in addition to any other options you use with the `MQOPEN` call. Note, however, that you cannot save context information if you only browse the message.

When you create the second message:

- Open the queue using the `MQOO_PASS_IDENTITY_CONTEXT` option (in addition to the `MQOO_OUTPUT` option).
- In the *Context* field of the put-message options structure, give the handle of the queue from which you saved the context information.
- In the *Options* field of the put-message options structure, specify the `MQPMO_PASS_IDENTITY_CONTEXT` option.

Passing user context

You cannot choose to pass only user context. To pass user context when putting a message, specify `MQPMO_PASS_ALL_CONTEXT`. Any properties in the user context are passed in the same way as the origin context.

When an `MQPUT` or `MQPUT1` takes place and the context is being passed, all properties in the user context are passed from the retrieved message to the put message. Any user context properties that the

putting application has altered are put with their original values. Any user context properties that the putting application has deleted are restored in the put message. Any user context properties that the putting application has added to the message are retained.

Passing all context

If your application gets a message, and puts the message data (unchanged) into another message, the application must pass all (identity, origin, and user) context information from the original message to the new message. An example of an application that might do this is a message mover, which moves messages from one queue to another.

Follow the same procedure as for passing identity context, except that you use the MQOPEN option MQOO_PASS_ALL_CONTEXT and the put-message option MQPMO_PASS_ALL_CONTEXT.

Setting identity context

If you want to set the identity context information for a message:

- Open the queue using the MQOO_SET_IDENTITY_CONTEXT option.
- Put the message on the queue, specifying the MQPMO_SET_IDENTITY_CONTEXT option. In the message descriptor, specify whatever identity context information you require.

Note: When you set some (but not all) of the identity context fields using the MQOO_SET_IDENTITY_CONTEXT and MQPMO_SET_IDENTITY_CONTEXT options, it is important to realize that the queue manager does not set any of the other fields.

In order to modify any of the message context options, you must have the appropriate authorizations to issue the call. For example, in order to use MQOO_SET_IDENTITY_CONTEXT or MQPMO_SET_IDENTITY_CONTEXT, you must have +setid permission.

Setting user context

To set a property in the user context, set the Context field of the message property descriptor (MQPD) to MQPD_USER_CONTEXT when you make the MQSETMP call.

You do not need any special authority to set a property in the user context. User context has no MQOO_SET_* or MQPMO_SET_* context options.

Setting all context

If you want to set both the identity and the origin context information for a message:

1. Open the queue using the MQOO_SET_ALL_CONTEXT option.
2. Put the message on the queue, specifying the MQPMO_SET_ALL_CONTEXT option. In the message descriptor, specify whatever identity and origin context information you require.

Appropriate authority is needed for each type of context setting.

Related concepts:

“Message context” on page 40

Message context information allows the application that retrieves the message to find out about the originator of the message.

Related reference:

“MQOPEN options relating to message context” on page 108

If you want to be able to associate context information with a message when you put it on a queue, you must use one of the message context options when you open the queue.

Putting one message on a queue using the MQPUT1 call:

Use the MQPUT1 call when you want to close the queue immediately after you have put a single message on it. For example, a server application is likely to use the MQPUT1 call when it is sending a reply to each of the different queues.

MQPUT1 is functionally equivalent to calling MQOPEN followed by MQPUT, followed by MQCLOSE. The only difference in the syntax for the MQPUT and MQPUT1 calls is that for MQPUT you specify an object handle, whereas for MQPUT1 you specify an object descriptor structure (MQOD) as defined in MQOPEN (see “Identifying objects (the MQOD structure)” on page 102). This is because you need to give information to the MQPUT1 call about the queue that it has to open, whereas when you call MQPUT, the queue must already be open.

As input to the MQPUT1 call, you must supply:

- A connection handle.
- A description of the object that you want to open. This is in the form of an object descriptor structure (MQOD).
- A description of the message that you want to put on the queue. This is in the form of a message descriptor structure (MQMD).
- Control information in the form of a put-message options structure (MQPMO).
- The length of the data contained within the message (MQLONG).
- The address of the message data.

The output from MQPUT1 is:

- A completion code
- A reason code

If the call completes successfully, it also returns your options structure and your message descriptor structure. The call modifies your options structure to show the name of the queue and the queue manager to which the message was sent. If you request that the queue manager generate a unique value for the identifier of the message that you are putting (by specifying binary zero in the *MsgId* field of the MQMD structure), the call inserts the value in the *MsgId* field before returning this structure to you.

Note: You cannot use MQPUT1 with a model queue name; however, once a model queue has been opened, you can issue an MQPUT1 to the dynamic queue.

The six input parameters for MQPUT1 are:

Hconn This is a connection handle. For CICS applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN or MQCONNX call. For other programs, always use the connection handle returned by the MQCONN or MQCONNX call.

ObjDesc

This is an object descriptor structure (MQOD).

In the *ObjectName* and *ObjectQMgrName* fields, give the name of the queue on which you want to put a message, and the name of the queue manager that owns this queue.

The *DynamicQName* field is ignored for the MQPUT1 call because it cannot use model queues.

Use the *AlternateUserId* field if you want to nominate an alternate user identifier that is to be used to test authority to open the queue.

MsgDesc

This is a message descriptor structure (MQMD). As with the MQPUT call, use this structure to define the message that you are putting on the queue.

PutMsgOpts

This is a put-message options structure (MQPMO). Use it as you would for the MQPUT call (see “Specifying options using the MQPMO structure” on page 113).

When the *Options* field is set to zero, the queue manager uses your own user ID when it performs tests for authority to access the queue. Also, the queue manager ignores any alternate user identifier given in the *AlternateUserId* field of the MQOD structure.

BufferLength

This is the length of your message.

Buffer This is the buffer area that contains the text of your message.

When you use clusters, MQPUT1 operates as though MQOO_BIND_NOT_FIXED is in effect. Applications must use the resolved fields in the MQPMO structure rather than the MQOD structure to determine where the message was sent. See Configuring a queue manager cluster for more information.

There is a description of the MQPUT1 call in MQPUT1.

Distribution lists:

Not supported on IBM MQ for z/OS. Distribution lists allow you to put a message to multiple destinations in a single MQPUT or MQPUT1 call.

A single MQOPEN call can open multiple queues and a single MQPUT call can then put a message to each of those queues. Some generic information from the MQI structures used for this process can be superseded by specific information relating to the individual destinations included in the distribution list.



Attention: Distribution lists do not support the use of alias queues that point to topic objects. From Version 9.0.1 and Version 9.0.0, Fix Pack 1, if an alias queue points to a topic object in a distribution list, IBM MQ returns MQRC_ALIAS_BASE_Q_TYPE_ERROR.

When an MQOPEN call is issued, generic information is taken from the Object Descriptor (MQOD). If you specify MQOD_VERSION_2 in the *Version* field and a value greater than zero in the *RecsPresent* field, the *Hobj* can be defined as a handle of a list (of one or more queues) rather than of a queue. In this case, specific information is given through the object records (MQORs), which give details of destination (that is, *ObjectName* and *ObjectQMgrName*).

The object handle (*Hobj*) is passed to the MQPUT call, allowing you to put to a list rather than to a single queue.

When a message is put on the queues (MQPUT), generic information is taken from the Put Message Option structure (MQPMO) and the Message Descriptor (MQMD). Specific information is given in the form of Put Message Records (MQPMRs).

Response Records (MQRR) can receive a completion code and reason code specific to each destination queue.

Figure 9 shows how distribution lists work.

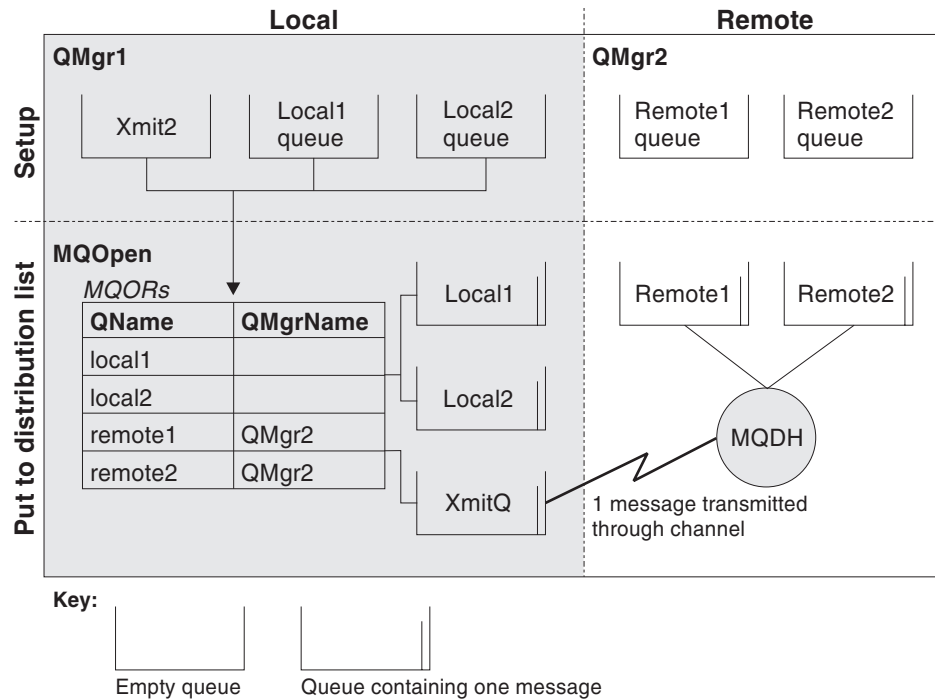


Figure 9. How distribution lists work. This diagram shows that *one* message is transmitted through the channel and can be put on more than one remote queue.

Opening distribution lists:

Use the MQOPEN call to open a distribution list, and use the options of the call to specify what you want to do with the list.

As input to MQOPEN, you must supply:

- A connection handle (see “Putting messages on a queue” on page 111 for a description)
- Generic information in the Object Descriptor structure (MQOD)
- The name of each queue that you want to open, using the Object Record structure (MQOR)

The output from MQOPEN is:

- An object handle that represents your access to the distribution list
- A generic completion code
- A generic reason code
- Response Records (optional), containing a completion code and reason for each destination

Using the MQOD structure

Use the MQOD structure to identify the queues that you want to open.

To define a distribution list, you must specify MQOD_VERSION_2 in the *Version* field, a value greater than zero in the *RecsPresent* field, and MQOT_Q in the *ObjectType* field. See MQOD for a description of all the fields of the MQOD structure.

Using the MQOR structure

Provide an MQOR structure for each destination.

The structure contains the destination queue and queue manager names. The *ObjectName* and *ObjectQMgrName* fields in the MQOD are not used for distribution lists. There must be one or more object records. If the *ObjectQMgrName* is left blank, the local queue manager is used. See *ObjectName* and *ObjectQMgrName* for further information about these fields.

You can specify the destination queues in two ways:

- By using the offset field *ObjectRecOffset*.

In this case, the application must declare its own structure containing an MQOD structure, followed by the array of MQOR records (with as many array elements as are needed), and set *ObjectRecOffset* to the offset of the first element in the array from the start of the MQOD. Ensure that this offset is correct. Use of built-in facilities provided by the programming language is recommended, if these are available in all the environments in which the application runs. The following code illustrates this technique for the COBOL programming language:

```
01 MY-OPEN-DATA.  
  02 MY-MQOD.  
    COPY CMQODV.  
  02 MY-MQOR-TABLE OCCURS 100 TIMES.  
    COPY CMQORV.  
  MOVE LENGTH OF MY-MQOD TO MQOD-OBJECTRECOFFSET.
```

Alternatively, use the constant MQOD_CURRENT_LENGTH if the programming language does not support the necessary built-in facilities in all the environments concerned. The following code illustrates this technique:

```
01 MY-MQ-CONSTANTS.  
  COPY CMQV.  
01 MY-OPEN-DATA.  
  02 MY-MQOD.  
    COPY CMQODV.  
  02 MY-MQOR-TABLE OCCURS 100 TIMES.  
    COPY CMQORV.  
  MOVE MQOD-CURRENT-LENGTH TO MQOD-OBJECTRECOFFSET.
```

However, this works correctly only if the MQOD structure and the array of MQOR records are contiguous; if the compiler inserts skip bytes between the MQOD and the MQOR array, these must be added to the value stored in *ObjectRecOffset*.

Using *ObjectRecOffset* is recommended for programming languages that do not support the pointer data type, or that implement the pointer data type in a way that is not portable to different environments (for example, the COBOL programming language).

- By using the pointer field *ObjectRecPtr*.

In this case, the application can declare the array of MQOR structures separately from the MQOD structure, and set *ObjectRecPtr* to the address of the array. The following code illustrates this technique for the C programming language:

```
MQOD MyMqod;  
MQOR MyMqor[100];  
MyMqod.ObjectRecPtr = MyMqor;
```

Using *ObjectRecPtr* is recommended for programming languages that support the pointer data type in a way that is portable to different environments (for example, the C programming language).

Whichever technique you choose, you must use one of *ObjectRecOffset* and *ObjectRecPtr*; the call fails with reason code MQRC_OBJECT_RECORDS_ERROR if both are zero, or both are nonzero.

Using the MQRR structure

These structures are destination-specific; each Response Record contains a *CompCode* and *Reason* field for each queue of a distribution list. You must use this structure to enable you to distinguish where any problems lie.

For example, if you receive a reason code of MQRC_MULTIPLE_REASONS and your distribution list contains five destination queues, you will not know which queues the problems apply to if you do not use this structure. However, if you have a completion code and reason code for each destination, you can locate the errors more easily.

See MQRR for further information about the MQRR structure.

Figure 10 shows how you can open a distribution list in C.




Figure 10. Opening a distribution list in C. The MQOD uses pointers to the MQOR and MQRR structures.

Figure 11 shows how you can open a distribution list in COBOL.




Figure 11. Opening a distribution list in COBOL. The MQOD uses offsets in COBOL.

Using the MQOPEN options

You can specify the following options when opening a distribution list:

- MQOO_OUTPUT
- MQOO_FAIL_IF QUIESCING (optional)
- MQOO_ALTERNATE_USER_AUTHORITY (optional)
- MQOO_*_CONTEXT (optional)

See “Opening and closing objects” on page 100 for a description of these options.

Putting messages to a distribution list:

To put messages to a distribution list, you can use MQPUT or MQPUT1.

As input, you must supply:

- A connection handle (see “Putting messages on a queue” on page 111 for a description).
- An object handle. If a distribution list is opened using MQOPEN, the *Hobj* allows you only to put to the list.
- A message descriptor structure (MQMD). See MQMD for a description of this structure.
- Control information in the form of a put-message option structure (MQPMO). See “Specifying options using the MQPMO structure” on page 113 for information about completing the fields of the MQPMO structure.
- Control information in the form of Put Message Records (MQPMR).
- The length of the data contained within the message (MQLONG).
- The message data itself.

The output is:

- A completion code
- A reason code
- Response Records (optional)

Using the MQPMR structure

This structure is optional and gives destination-specific information for some fields that you might want to identify differently from those already identified in the MQMD.

For a description of these fields, see MQPMR.

The content of each record depends on the information given in the *PutMsgRecFields* field of the MQPMO. For example, in the sample program AMQSPTL0.C (see “The Distribution List sample program” on page 490 for a description) showing the use of distribution lists, the sample chooses to provide values for *MsgId* and *CorrelId* in the MQPMR. This section of the sample program looks like this:

```
typedef struct
{
MQBYTE24 MsgId;
MQBYTE24 CorrelId;
} PutMsgRec;
...
/*****
MQLONG PutMsgRecFields=MQPMRF_MSG_ID | MQPMRF_CORREL_ID;
```

This implies that *MsgId* and *CorrelId* are provided for each destination of a distribution list. The Put Message Records are provided as an array.

Figure 12 on page 127 shows how you can put a message to a distribution list in C.




Figure 12. Putting a message to a distribution list in C. The MQPMO uses pointers to the MQPMR and MQRR structures.

Figure 13 shows how you can put a message to a distribution list in COBOL.




Figure 13. Putting a message to a distribution list in COBOL. The MQPMO uses offsets in COBOL.

Using MQPUT1

If you are using MQPUT1, consider the following points:

1. The values of the *ResponseRecOffset* and *ResponseRecPtr* fields must be null or zero.
2. The Response Records, if required, must be addressed from the MQOD.

Some cases where the put calls fail:

If certain attributes of a queue are changed using the FORCE option on a command during the interval between you issuing an MQOPEN and an MQPUT call, the MQPUT call fails and returns the MQRC_OBJECT_CHANGED reason code.

The queue manager marks the object handle as being no longer valid. This also happens if the changes are made while an MQPUT1 call is being processed, or if the changes apply to any queue to which the queue name resolves. The attributes that affect the handle in this way are listed in the description of the MQOPEN call in MQOPEN. If your call returns the MQRC_OBJECT_CHANGED reason code, close the queue, reopen it, then try to put a message again.

If put operations are inhibited for a queue on which you are attempting to put messages (or any queue to which the queue name resolves), the MQPUT or MQPUT1 call fails and returns the MQRC_PUT_INHIBITED reason code. You might be able to put a message successfully if you attempt the call at a later time, if the design of the application is such that other programs change the attributes of queues regularly.

Furthermore, if the queue that you are trying to put your message on is full, the MQPUT or MQPUT1 call fails and returns MQRC_Q_FULL.

If a dynamic queue (either temporary or permanent) has been deleted, MQPUT calls using a previously-acquired object handle fail and return the MQRC_Q_DELETED reason code. In this situation, it is good practice to close the object handle as it is no longer of any use to you.

In the case of distribution lists, multiple completion codes and reason codes can occur in a single request. These cannot be handled using only the *CompCode* and *Reason* output fields on MQOPEN and MQPUT.

When you use distribution lists to put messages to multiple destinations, the Response Records contain the specific *CompCode* and *Reason* for each destination. If you receive a completion code of MQCC_FAILED, no message is put on any destination queue successfully. If the completion code is MQCC_WARNING, the message is successfully put on one or more of the destination queues. If you receive a return code of MQRC_MULTIPLE_REASONS, the reason codes are not all the same for every destination. Therefore, it is recommended to use the MQRR structure so that you can determine which queue or queues caused an error and the reasons for each.

Getting messages from a queue

Use this information to learn about getting messages from a queue.



You can get messages from a queue in two ways:

1. You can remove a message from the queue so that other programs can no longer see it.
2. You can copy a message, leaving the original message on the queue. This is known as *browsing*. You can remove the message once you have browsed it.

In both cases, you use the MQGET call, but first your application must be connected to the queue manager, and you must use the MQOPEN call to open the queue (for input, browse, or both). These operations are described in “Connecting to and disconnecting from a queue manager” on page 91 and “Opening and closing objects” on page 100.

When you have opened the queue, you can use the MQGET call repeatedly to browse or remove messages on the same queue. Call MQCLOSE when you have finished getting all the messages that you want from the queue.

Use the following links to find out more about getting messages from a queue:

- “Getting messages from a queue using the MQGET call” on page 129
- “The order in which messages are retrieved from a queue” on page 133
- “Getting a particular message” on page 146
- “Improving performance of non-persistent messages” on page 147
-  “Type of index” on page 152
- “Handling messages greater than 4 MB long” on page 152
- “Waiting for messages” on page 159
-  “Signaling” on page 160
- “Skipping backout” on page 161
- “Application data conversion” on page 164
- “Browsing messages on a queue” on page 165
- “Some cases where the MQGET call fails” on page 172

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

▶ **z/OS** “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

▶ **z/OS** “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

This information helps you to write IMS applications using IBM MQ.

Getting messages from a queue using the MQGET call:

The MQGET call gets a message from an open local queue. It cannot get a message from a queue on another system.

As input to the MQGET call, you must supply:

- A connection handle.
- A queue handle.
- A description of the message that you want to get from the queue. This is in the form of a message descriptor (MQMD) structure.
- Control information in the form of a Get Message Options (MQGMO) structure.
- The size of the buffer that you have assigned to hold the message (MQLONG).
- The address of the storage in which to put the message.

The output from MQGET is:


- A reason code
- A completion code
- The message in the buffer area that you specified, if the call completes successfully
- Your options structure, modified to show the name of the queue from which the message was retrieved
- Your message descriptor structure, with the contents of the fields modified to describe the message that was retrieved
- The length of the message (MQLONG)

There is a description of the MQGET call in MQGET.

The following sections describe the information you must supply as input to the MQGET call.

- “Specifying connection handles”
- “Describing messages using the MQMD structure and the MQGET call”
- “Specifying MQGET options using the MQGMO structure”
- “Specifying the size of the buffer area” on page 132

Specifying connection handles

 For CICS on z/OS applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN or MQCONNX call. For other applications, always use the connection handle returned by the MQCONN or MQCONNX call.

Use the queue handle (*Hobj*) that is returned when you call MQOPEN.

Describing messages using the MQMD structure and the MQGET call

To identify the message that you want to get from a queue, use the message descriptor structure (MQMD).

This is an input/output parameter for the MQGET call. There is an introduction to the message properties that MQMD describes in “IBM MQ messages” on page 9, and there is a description of the structure itself in MQMD.

If you know which message you want to get from the queue, see “Getting a particular message” on page 146.

If you do not specify a particular message, MQGET retrieves the *first* message in the queue. “The order in which messages are retrieved from a queue” on page 133 describes how the priority of a message, the **MsgDeliverySequence** attribute of the queue, and the MQGMO_LOGICAL_ORDER option determine the order of the messages in the queue.

Note: If you want to use MQGET more than once (for example, to step through the messages in the queue), you must set the *MsgId* and *CorrelId* fields of this structure to null after each call. This clears these fields of the identifiers of the message that was retrieved.

However, if you want to group your messages, the *GroupId* must be the same for messages in the same group, so that the call looks for a message having the same identifiers as the previous message in order to make up the whole group.

Specifying MQGET options using the MQGMO structure

The MQGMO structure is an input/output variable for passing options to the MQGET call. The following sections help you to complete some of the fields of this structure.

There is a description of the MQGMO structure in MQGMO.

StrucId

StrucId is a 4-character field used to identify the structure as a get-message options structure. Always specify MQGMO_STRUC_ID.







Version

Version describes the version number of the structure. MQGMO_VERSION_1 is the default. If you want to use the Version 2 fields or retrieve messages in logical order, specify MQGMO_VERSION_2. If you want to use the Version 3 fields or retrieve messages in logical order, specify MQGMO_VERSION_3. MQGMO_CURRENT_VERSION sets your application to use the most recent level.


Options


Within your code, you can select the options in any order; each option is represented by a bit in the *Options* field.


The *Options* field controls:

- Whether the MQGET call waits for a message to arrive on the queue before it completes (see “Waiting for messages” on page 159)
- Whether the get operation is included in a unit of work.
- Whether a nonpersistent message is retrieved outside sync point, allowing fast messaging
-  On IBM MQ for z/OS, whether the message retrieved is marked as skipping backout (see “Skipping backout” on page 161)
- Whether the message is removed from the queue, or merely browsed
- Whether to select a message by using a browse cursor or by other selection criteria
- Whether the call succeeds even if the message is longer than your buffer
-  On IBM MQ for z/OS, whether to allow the call to complete. This option also sets a signal to indicate that you want to be notified when a message arrives
- Whether the call fails if the queue manager is in a quiescing state
-  On IBM MQ for z/OS, whether the call fails if the connection is in a quiescing state
- Whether application message data conversion is required (see “Application data conversion” on page 164)
- The order in which messages and segments are retrieved from a queue  (except for IBM MQ for z/OS)
- Whether complete, logical messages only are retrievable  (except for IBM MQ for z/OS)
- Whether messages in a group can be retrieved only when *all* messages in the group are available
- Whether segments in a logical message can be retrieved only when *all* segments in the logical message are available  (except for IBM MQ for z/OS)

If you leave the *Options* field set to the default value (MQGMO_NO_WAIT), the MQGET call operates this way:

- If there is no message matching your selection criteria on the queue, the call does not wait for a message to arrive, but completes immediately.  Also, in IBM MQ for z/OS, the call does not set a signal requesting notification when such a message arrives.
- The way that the call operates with sync points is determined by the platform:

Platform	Under sync point control
IBM i	No
UNIX and Linux systems	No
 z/OS	Yes
Windows systems	No

-  On IBM MQ for z/OS, the message retrieved is not marked as skipping backout.
- The selected message is removed from the queue (not browsed).
- No application message data conversion is required.
- The call fails if the message is longer than your buffer.

WaitInterval

The *WaitInterval* field specifies the maximum time (in milliseconds) that the MQGET call waits

for a message to arrive on the queue when you use the MQGMO_WAIT option. If no message arrives within the time specified in *WaitInterval*, the call completes and returns a reason code showing that there was no message that matched your selection criteria on the queue.

► **z/OS** On IBM MQ for z/OS, if you use the MQGMO_SET_SIGNAL option, the *WaitInterval* field specifies the time for which the signal is set.

For more information about these options, see “Waiting for messages” on page 159 ► **z/OS** and “Signaling” on page 160.

Signal1

Signal1 is supported only on ► **z/OS** IBM MQ for z/OS and MQSeries® for HP Integrity NonStop Server.

If you use the MQGMO_SET_SIGNAL option to request that your application is notified when a suitable message arrives, you specify the type of signal in the *Signal1* field. In IBM MQ on all other platforms, the *Signal1* field is reserved and its value is not significant.

► **z/OS** For more information, see “Signaling” on page 160.

Signal2

The *Signal2* field is reserved on all platforms and its value is not significant.

► **z/OS** For more information, see “Signaling” on page 160.

ResolvedQName

ResolvedQName is an output field in which the queue manager returns the name of the queue (after resolution of any alias) from which the message was retrieved.

MatchOptions

MatchOptions controls the selection criteria for MQGET.

GroupStatus

GroupStatus indicates whether the message that you have retrieved is in a group.

SegmentStatus

SegmentStatus indicates whether the item that you have retrieved is a segment of a logical message.

Segmentation

Segmentation indicates whether segmentation is allowed for the message retrieved.

MsgToken

MsgToken uniquely identifies a message.

ReturnedLength

ReturnedLength is an output field in which the queue manager returns the length of message data returned (in bytes).

MsgHandle

The handle to a message that is to be populated with the properties of the message being retrieved from the queue. The handle has previously been created by an MQCRTMH call. Any properties already associated with the handle are cleared before retrieving a message.

Specifying the size of the buffer area

In the **BufferLength** parameter of the MQGET call, specify the size of the buffer area to hold the message data that you retrieve. You decide how large this should be in three ways:

1. You might already know what length of messages to expect from this program. If so, specify a buffer of this size.

However, you can use the `MQGMO_ACCEPT_TRUNCATED_MSG` option in the `MQGMO` structure if you want the `MQGET` call to complete even if the message is too large for the buffer. In this case:

- The buffer is filled with as much of the message as it can hold
- The call returns a warning completion code
- The message is removed from the queue (discarding the remainder of the message), or the browse cursor is advanced (if you are browsing the queue)
- The real length of the message is returned in *DataLength*

Without this option, the call still completes with a warning, but it does not remove the message from the queue (or advance the browse cursor).

2. Estimate a size for the buffer (or even specify a size of zero bytes) and *do not* use the `MQGMO_ACCEPT_TRUNCATED_MSG` option. If the `MQGET` call fails (for example, because the buffer is too small), the length of the message is returned in the **DataLength** parameter of the call. (The buffer is still filled with as much of the message as it can hold, but the processing of the call is not completed.) Store the *MsgId* of this message, then repeat the `MQGET` call, specifying a buffer area of the correct size, and the *MsgId* that you noted from the first call.

If your program is serving a queue that is also being served by other programs, one of those other programs might remove the message that you want before your program can issue another `MQGET` call. Your program could waste time searching for a message that no longer exists. To avoid this, first browse the queue until you find the message that you want, specifying a *BufferLength* of zero and using the `MQGMO_ACCEPT_TRUNCATED_MSG` option. This positions the browse cursor under the message that you want. You can then retrieve the message by calling `MQGET` again, specifying the `MQGMO_MSG_UNDER_CURSOR` option. If another program removes the message between your browse and removal calls, your second `MQGET` fails immediately (without searching the whole queue), because there is no message under your browse cursor.

3. The *MaxMsgLength queue* attribute determines the maximum length of messages accepted for that queue; the *MaxMsgLength queue manager* attribute determines the maximum length of messages accepted for that queue manager. If you do not know what length of message to expect, you can inquire about the **MaxMsgLength** attribute (using the `MQINQ` call), then specify a buffer of this size.

Try to make the buffer size as close as possible to the actual message size to avoid reduced performance.

For further information about the **MaxMsgLength** attribute, see “Increasing the maximum message length” on page 153.

The order in which messages are retrieved from a queue:

You can control the order in which you retrieve messages from a queue. This section looks at the options.

Priority:

A program can assign a priority to a message when it puts the message on a queue (see “Message priorities” on page 17). Messages of equal priority are stored in a queue in order of arrival, not the order in which they are committed.


The queue manager maintains queues either in strict FIFO (first in, first out) sequence, or in FIFO within priority sequence. This depends on the setting of the **MsgDeliverySequence** attribute of the queue. When a message arrives on a queue, it is inserted immediately following the last message that has the same priority.

Programs can either get the first message from a queue, or they can get a particular message from a queue, ignoring the priority of those messages. For example, a program might want to process the reply to a particular message that it sent earlier. For more information, see “Getting a particular message” on page 146.

If an application puts a sequence of messages on a queue, another application can retrieve those messages in the same order that they were put, provided:

- The messages all have the same priority
- The messages were all put within the same unit of work, or all put outside a unit of work
- The queue is local to the putting application

If these conditions are not met, and the applications depend on the messages being retrieved in a certain order, the applications must either include sequencing information in the message data, or establish a means of acknowledging receipt of a message before the next one is sent.

 On IBM MQ for z/OS, you can use the queue attribute, *IndexType*, to increase the speed of MQGET operations on the queue. For more information, see “Type of index” on page 152.

Logical and physical ordering:

Messages on queues can occur (within each priority level) in *physical* or *logical* order.

Physical order is the order in which messages arrive on a queue. Logical order is when all of the messages and segments within a group are in their logical sequence, next to each other, in the position determined by the physical position of the first item belonging to the group.

For a description of groups, messages, and segments, see “Message groups” on page 36. These physical and logical orders can differ because:

- Groups can arrive at a destination at similar times from different applications, therefore losing any distinct physical order.
- Even within a single group, messages can get out of order because of rerouting or delay of some of the messages in the group.

For example, the logical order might look like Figure Figure 14 on page 135:




Figure 14. Logical order on a queue

These messages would occur in the following logical order on a queue:

1. Message A (not in a group)
2. Logical message 1 of group Y
3. Logical message 2 of group Y
4. Segment 1 of (last) logical message 3 of group Y
5. (Last) segment 2 of (last) logical message 3 of group Y
6. Logical message 1 of group Z
7. (Last) logical message 2 of group Z
8. Message B (not in a group)

The physical order, however, might be entirely different. The physical position of the *first* item within each group determines the logical position of the whole group. For example, if groups Y and Z arrived at similar times, and message 2 of group Z overtook message 1 of the same group, the physical order would look like Figure Figure 15 on page 136:

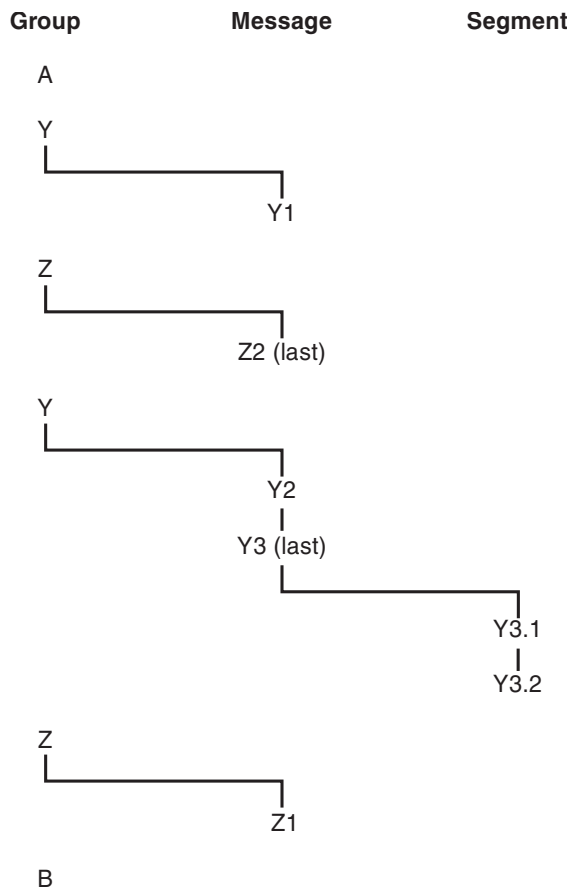


Figure 15. Physical order on a queue

These messages occur in the following physical order on the queue:

1. Message A (not in a group)
2. Logical message 1 of group Y
3. Logical message 2 of group Z
4. Logical message 2 of group Y
5. Segment 1 of (last) logical message 3 of group Y
6. (Last) segment 2 of (last) logical message 3 of group Y
7. Logical message 1 of group Z
8. Message B (not in a group)

Note: On IBM MQ for z/OS, the physical order of messages on the queue is not guaranteed if the queue is indexed by GROUPID.

When getting messages, you can specify MQGMO_LOGICAL_ORDER to retrieve messages in logical order rather than physical order.

If you issue an MQGET call with MQGMO_BROWSE_FIRST and MQGMO_LOGICAL_ORDER, subsequent MQGET calls with MQGMO_BROWSE_NEXT must also specify MQGMO_LOGICAL_ORDER. Conversely, if the MQGET with MQGMO_BROWSE_FIRST does not specify MQGMO_LOGICAL_ORDER, neither must the following MQGETs with MQGMO_BROWSE_NEXT.

The group and segment information that the queue manager retains for MQGET calls that browse messages on the queue is separate from the group and segment information that the queue manager retains for MQGET calls that remove messages from the queue. When you specify MQGMO_BROWSE_FIRST, the queue manager ignores the group and segment information for browsing, and scans the queue as though there were no current group and no current logical message.

Note: Do not use an MQGET call to browse *beyond the end* of a message group (or logical message not in a group) without specifying MQGMO_LOGICAL_ORDER. For example, if the last message in the group precedes the first message in the group on the queue, using MQGMO_BROWSE_NEXT to browse beyond the end of the group, specifying MQMO_MATCH_MSG_SEQ_NUMBER with *MsgSeqNumber* set to 1 (to find the first message of the next group) returns again the first message in the group already browsed. This could happen immediately, or a number of MQGET calls later (if there are intervening groups).

Avoid the possibility of an infinite loop by opening the queue *twice* for browse:

- Use the first handle to browse only the first message in each group.
- Use the second handle to browse only the messages within a specific group.
- Use the MQMO_* options to move the second browse cursor to the position of the first browse cursor, before browsing the messages in the group.
- Do not use the MQGMO_BROWSE_NEXT browse beyond the end of a group.

For further information about this, see MQGET, MQMD, and Rules for validating MQI options.

For most applications you will probably choose either logical or physical ordering when browsing. However, if you want to switch between these modes, remember that when you first issue a browse with MQGMO_LOGICAL_ORDER, your position within the logical sequence is established.

If the first item within the group is not present at this time, the group that you are in is not considered to be part of the logical sequence.

Once the browse cursor is within a group, it can continue within the same group, even if the first message is removed. Initially though, you can never move into a group using MQGMO_LOGICAL_ORDER where the first item is not present.

MQPMO_LOGICAL_ORDER

The MQPMO option tells the queue manager how the application puts messages in groups and segments of logical messages. It can be specified only on the MQPUT call; it is not valid on the MQPUT1 call.

If MQPMO_LOGICAL_ORDER is specified, it indicates that the application uses successive MQPUT calls to:

1. Put the segments in each logical message in the order of increasing segment offset, starting from 0, with no gaps.
2. Put all the segments in one logical message before putting the segments in the next logical message.
3. Put the logical messages in each message group in the order of increasing message sequence number, starting from 1, with no gaps. IBM MQ increments the message sequence number automatically.
4. Put all the logical messages in one message group before putting logical messages in the next message group.

Because the application has told the queue manager how it puts messages in groups and segments of logical messages, the application does not have to maintain and update the group and segment information about each MQPUT call, because the queue manager maintains and updates this information. Specifically, it means that the application does not need to set the *GroupId*, *MsgSeqNumber*, and *Offset* fields in MQMD, because the queue manager sets these fields to the appropriate values. The application must only set the *MsgFlags* field in MQMD, to indicate

when messages belong to groups or are segments of logical messages, and to indicate the last message in a group or last segment of a logical message.

After a message group or logical message has been started, subsequent MQPUT calls must specify the appropriate MQMF_* flags in *MsgFlags* in MQMD. If the application tries to put a message that is not in a group when there is an unterminated message group, or put a message that is not a segment when there is an unterminated logical message, the call fails with reason code MQRC_INCOMPLETE_GROUP or MQRC_INCOMPLETE_MSG, as appropriate. However, the queue manager retains the information about the current message group or current logical message, and the application can terminate them by sending a message (possibly with no application message data) specifying MQMF_LAST_MSG_IN_GROUP or MQMF_LAST_SEGMENT as appropriate, before reissuing the MQPUT call to put the message that is not in the group or not a segment.

Figure 15 on page 136 shows the combinations of options and flags that are valid, and the values of the *GroupId*, *MsgSeqNumber*, and *Offset* fields that the queue manager uses in each case. Combinations of options and flags that are not shown in the table are not valid. The columns in the table have the following meanings; Either means Yes or No:

LOG ORD

Whether the MQPMO_LOGICAL_ORDER option is specified on the call.

MIG Whether the MQMF_MSG_IN_GROUP or MQMF_LAST_MSG_IN_GROUP option is specified on the call.

SEG Whether the MQMF_SEGMENT or MQMF_LAST_SEGMENT option is specified on the call.

SEG OK

Whether the MQMF_SEGMENTATION_ALLOWED option is specified on the call.

Cur grp

Whether a current message group exists before the call.

Cur log msg

Whether a current logical message exists before the call.

Other columns

Show the values that the queue manager uses. Previous denotes the value used for the field in the previous message for the queue handle.

Table 18. MQPUT options relating to messages in groups and segments of logical messages

Options you specify	Options you specify	Options you specify	Options you specify	Group and log-msg status before call	Group and log-msg status before call	Values the queue manager uses	Values the queue manager uses	Values the queue manager uses
LOG ORD	MIG	SEG	SEG OK	Cur grp	Cur log msg	<i>GroupId</i>	<i>MsgSeqNumber</i>	<i>Offset</i>
Yes	No	No	No	No	No	MQGL_NONE	1	0
Yes	No	No	Yes	No	No	New group id	1	0
Yes	No	Yes	Either	No	No	New group id	1	0
Yes	No	Yes	Either	No	Yes	Previous group id	1	Previous offset + previous segment length
Yes	Yes	Either	Either	No	No	New group id	1	0

Table 18. MQPUT options relating to messages in groups and segments of logical messages (continued)

Options you specify	Options you specify	Options you specify	Options you specify	Group and log-msg status before call	Group and log-msg status before call	Values the queue manager uses	Values the queue manager uses	Values the queue manager uses
Yes	Yes	Either	Either	Yes	No	Previous group id	Previous sequence number + 1	0
Yes	Yes	Yes	Either	Yes	Yes	Previous group id	Previous sequence number	Previous offset + previous segment length
No	No	No	No	Either	Either	MQGL_NONE	1	0
No	No	No	Yes	Either	Either	New group ID if MQGL_NONE, else value in field	1	0
No	No	Yes	Either	Either	Either	New group ID if MQGL_NONE, else value in field	1	Value in field
No	Yes	No	Either	Either	Either	New group ID if MQGL_NONE, else value in field	Value in field	0
No	Yes	Yes	Either	Either	Either	New group ID if MQGL_NONE, else value in field	Value in field	Value in field

Note:

- MQPMO_LOGICAL_ORDER is not valid on the MQPUT1 call.
- For the *MsgId* field, the queue manager generates a new message identifier if MQPMO_NEW_MSG_ID or MQMI_NONE is specified, and uses the value in the field otherwise.
- For the *CorrelId* field, the queue manager generates a new correlation identifier if MQPMO_NEW_CORREL_ID is specified, and uses the value in the field otherwise.

When you specify MQPMO_LOGICAL_ORDER, the queue manager requires that all messages in a group and segments in a logical message are put with the same value in the *Persistence* field in MQMD, that is, all must be persistent, or all must be nonpersistent. If this condition is not satisfied, the MQPUT call fails with reason code MQRC_INCONSISTENT_PERSISTENCE.

The MQPMO_LOGICAL_ORDER option affects units of work as follows:

- If the first physical message in a group or logical message is put within a unit of work, all the other physical messages in the group or logical message must be put within a unit of work, if the same queue handle is used. However, they do not need to be put within the same unit of work, allowing a message group or logical message that consists of many physical messages to be split across two or more consecutive units of work for the queue handle.
- If the first physical message in a group or logical message is not put within a unit of work, none of the other physical messages in the group or logical message can be put within a unit of work, if the same queue handle is used.

If these conditions are not satisfied, the MQPUT call fails with reason code MQRC_INCONSISTENT_UOW.

When MQPMO_LOGICAL_ORDER is specified, the MQMD supplied on the MQPUT call must not be less than MQMD_VERSION_2. If this condition is not satisfied, the call fails with reason code MQRC_WRONG_MD_VERSION.

If MQPMO_LOGICAL_ORDER is not specified, messages in groups and segments of logical messages can be put in any order, and it is not necessary to put complete message groups or complete logical messages. It is the responsibility of the application to ensure that the *GroupId*, *MsgSeqNumber*, *Offset*, and *MsgFlags* fields have appropriate values.

Use this technique to restart a message group or logical message in the middle, after a system failure has occurred. When the system restarts, the application can set the *GroupId*, *MsgSeqNumber*, *Offset*, *MsgFlags*, and *Persistence* fields to the appropriate values, and then issue the MQPUT call with MQPMO_SYNCPOINT or MQPMO_NO_SYNCPOINT set as required, but without specifying MQPMO_LOGICAL_ORDER. If this call is successful, the queue manager retains the group and segment information, and subsequent MQPUT calls using that queue handle can specify MQPMO_LOGICAL_ORDER as normal.

The group and segment information that the queue manager retains for the MQPUT call is separate from the group and segment information that it retains for the MQGET call.

For any given queue handle, the application can mix MQPUT calls which specify MQPMO_LOGICAL_ORDER with MQPUT calls that do not, but note the following points:

- If MQPMO_LOGICAL_ORDER is not specified, each successful MQPUT call causes the queue manager to set the group and segment information for the queue handle to the values specified by the application, replacing the existing group and segment information retained by the queue manager for the queue handle.
- If MQPMO_LOGICAL_ORDER is not specified, the call does not fail if there is a current message group or logical message; the call might succeed with an MQCC_WARNING completion code. Table 19 shows the various cases that can arise. In these cases, if the completion code is not MQCC_OK, the reason code is one of the following (as appropriate):
 - MQRC_INCOMPLETE_GROUP
 - MQRC_INCOMPLETE_MSG
 - MQRC_INCONSISTENT_PERSISTENCE
 - MQRC_INCONSISTENT_UOW

Note: The queue manager does not check the group and segment information for the MQPUT1 call.

Table 19. Outcome when MQPUT or MQCLOSE call is not consistent with group and segment information

Current call is	Previous call was MQPUT with MQPMO_LOGICAL_ORDER	Previous call was MQPUT without MQPMO_LOGICAL_ORDER
MQPUT with MQPMO_LOGICAL_ORDER	MQCC_FAILED	MQCC_FAILED
MQPUT without MQPMO_LOGICAL_ORDER	MQCC_WARNING	MQCC_OK
MQCLOSE with an unterminated group or logical message	MQCC_WARNING	MQCC_OK

For applications that put messages and segments in logical order, specify MQPMO_LOGICAL_ORDER, as it is the simplest option to use. This option relieves the application of the need to manage the group and segment information, because the queue manager manages that information. However, specialized applications might need more control than that provided by the MQPMO_LOGICAL_ORDER option, which can be achieved by not specifying that option; if you do so, you must ensure that the *GroupId*, *MsgSeqNumber*, *Offset*, and *MsgFlags* fields in MQMD are set correctly, before each MQPUT or MQPUT1 call.

For example, an application that wants to forward physical messages that it receives, without regard for whether those messages are in groups or segments of logical messages, must not specify MQPMO_LOGICAL_ORDER, for two reasons:

- If the messages are retrieved and put in order, specifying MQPMO_LOGICAL_ORDER assigns a new group identifier to the messages, which might make it difficult or impossible for the originator of the messages to correlate any reply or report messages that result from the message group.
- In a complex network with multiple paths between sending and receiving queue managers, the physical messages might arrive out of order. By not specifying MQPMO_LOGICAL_ORDER and MQGMO_LOGICAL_ORDER on the MQGET call, the forwarding application can retrieve and forward each physical message as soon as it arrives, without waiting for the next one in logical order to arrive.

Applications that generate report messages for messages in groups or segments of logical messages must also not specify MQPMO_LOGICAL_ORDER when putting the report message.

MQPMO_LOGICAL_ORDER can be specified with any of the other MQPMO_* options.

Putting Logically Ordered Groups to a Clustered Queue (MQOO_BIND_ON_GROUP)

The MQOO_BIND_ON_OPEN option ensures that all messages from this application, and therefore all groups, are routed to a single instance. This has the drawback that the application traffic is not load balanced across multiple instances of a cluster queue. In order to enable workload balancing while keeping groups of messages intact, you must set the following options:

- The MQPUT call must specify MQPMO_LOGICAL_ORDER
- The MQOPEN call must specify one of the following two options:
 - MQOO_BIND_ON_GROUP
 - MQOO_BIND_AS_Q_DEF, and the queue definition must specify DEFBIND(GROUP)

Workload balancing is then driven *between groups* of messages without requiring an MQCLOSE and MQOPEN of the queue. *Between groups* means that MQMF_MSG_IN_GROUP is set in the MQMD(v2) or MQMDE, and there is no partially complete group in progress. When a group is in progress, the resolved queue manager and queue name in the object handle are reused.

If the previous message was MQPMO_LOGICAL_ORDER and/or MQMF_MSG_IN_GROUP was set but the current message is not part of the group, then the PUT call fails with MQRC_INCOMPLETE_GROUP.

If an individual MQPUT does not specify MQPMO_LOGICAL_ORDER, and no current group is active, then workload balancing is driven for that message (as if the MQOPEN call has specified MQOO_BIND_NOT_FIXED).

No reallocation takes place for messages bound to a destination using MQOO_BIND_ON_GROUP. For more information on reallocation, see “Message groups” on page 36.

Grouping logical messages:

There are two main reasons for using logical messages in a group:

- You might need to process the messages in a particular order.
- You might need to process each message in a group in a related way.

In either case, retrieve the entire group with the same getting application instance.

For example, assume that the group consists of four logical messages. The putting application looks like this:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP

MQCMIT
```

The getting application specifies the MQGMO_ALL_MSGS_AVAILABLE option for the first message in the group. This ensures that processing does not start until all the messages within the group have arrived. The MQGMO_ALL_MSGS_AVAILABLE option is ignored for subsequent messages within the group.

When the first logical message of the group is retrieved, you can use MQGMO_LOGICAL_ORDER to ensure that the remaining logical messages of the group are retrieved in order.

So, the getting application looks like this:

```
/* Wait for the first message in a group, or a message not in a group */
GMO.Options = MQGMO_SYNCPOINT | MQGMO_WAIT
             | MQGMO_ALL_MSGS_AVAILABLE | MQGMO_LOGICAL_ORDER
do while ( GroupStatus == MQGS_MSG_IN_GROUP )
  MQGET
  /* Process each remaining message in the group */
  ...

MQCMIT
```

For further examples of grouping messages, see “Application segmentation of logical messages” on page 156 and “Putting and getting a group that spans units of work” on page 143.

For information on allowing an application to request that a group of messages are all allocated to the same destination instance for cluster queues, see DefBind.

Putting and getting a group that spans units of work:

In the previous case, messages or segments cannot start to leave the node (if its destination is remote) or start to be retrieved until the whole group has been put and the unit of work is committed. This might not be what you want if it takes a long time to put the whole group, or if queue space is limited on the node. To overcome this, put the group in several units of work.

If the group is put within multiple units of work, it is possible for some of the group to commit even when the putting application fails. The application must therefore save status information, committed with each unit of work, which it can use after a restart to resume an incomplete group. The simplest place to record this information is in a STATUS queue. If a complete group has been successfully put, the STATUS queue is empty.

If segmentation is involved, the logic is similar. In this case, the **StatusInfo** must include the *Offset*.

Here is an example of putting the group in several units of work:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

/* First UOW */

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
StatusInfo = GroupId,MsgSeqNumber from MQMD
MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
MQCMIT

/* Next and subsequent UOWs */
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
StatusInfo = GroupId,MsgSeqNumber from MQMD
MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
MQCMIT

/* Last UOW */
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP
MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
MQCMIT
```

If all the units of work have been committed, the entire group has been put successfully, and the STATUS queue is empty. If not, the group must be resumed at the point indicated by the status information. MQPMO_LOGICAL_ORDER cannot be used for the first put, but can thereafter.

Restart processing looks like this:

```
MQGET (StatusInfo from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
if (Reason == MQRC_NO_MSG_AVAILABLE)
  /* Proceed to normal processing */
  ...
else
  /* Group was terminated prematurely */
  Set GroupId, MsgSeqNumber in MQMD to values from Status message
  PMO.Options = MQPMO_SYNCPOINT
  MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP

  /* Now normal processing is resumed.
  Assume this is not the last message */
  PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT
```

```

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
StatusInfo = GroupId,MsgSeqNumber from MQMD
MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
MQCMIT

```

From the getting application, you might want to start processing the messages in a group before the whole group has arrived. This improves response times on the messages within the group, and also means that storage is not required for the entire group. In order to realize the benefits, use several units of work for each group of messages. For recovery reasons, you must retrieve each message within a unit of work.

As with the corresponding putting application, this requires status information to be recorded somewhere automatically as each unit of work is committed. Again, the simplest place to record this information is on a STATUS queue. If a complete group has been successfully processed, the STATUS queue is empty.

Note: For intermediate units of work, you can avoid the MQGET calls from the STATUS queue by specifying that each MQPUT to the status queue is a segment of a message (that is, by setting the MQMF_SEGMENT flag), instead of putting a complete new message for each unit of work. In the last unit of work, a final segment is put to the status queue specifying MQMF_LAST_SEGMENT, and then the status information is cleared with an MQGET specifying MQGMO_COMPLETE_MSG.

During restart processing, instead of using a single MQGET to get a possible status message, browse the status queue with MQGMO_LOGICAL_ORDER until you reach the last segment (that is, until no further segments are returned). In the first unit of work after restart, also specify the offset explicitly when putting the status segment.

In the following example, we consider only messages within a group, assuming that the application's buffer is always large enough to hold the entire message, whether or not the message has been segmented. MQGMO_COMPLETE_MSG is therefore specified on each MQGET. The same principles apply if segmentation is involved (in this case, the StatusInfo must include the *offset*).

For simplicity, we assume that a maximum of 4 messages are retrieved within a single UOW:

```

msgs = 0    /* Counts messages retrieved within UOW */
/* Should be no status message at this point */

/* Retrieve remaining messages in the group */
do while ( GroupStatus == MQGS_MSG_IN_GROUP )

    /* Process up to 4 messages in the group */
    GMO.Options = MQGMO_SYNCPOINT | MQGMO_WAIT
                 | MQGMO_LOGICAL_ORDER
    do while ( (GroupStatus == MQGS_MSG_IN_GROUP) && (msgs < 4) )
        MQGET
        msgs = msgs + 1
        /* Process this message */
        ...
    /* end while

    /* Have retrieved last message or 4 messages */
    /* Update status message if not last in group */
    MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
    if ( GroupStatus == MQGS_MSG_IN_GROUP )
        StatusInfo = GroupId,MsgSeqNumber from MQMD
        MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
    MQCMIT
    msgs = 0
/* end while

```



```

if ( msgs > 0 )
  /* Come here if there was only 1 message in the group */
  MQCMIT

```

If all the units of work have been committed, the entire group has been retrieved successfully, and the STATUS queue is empty. If not, the group must be resumed at the point indicated by the status information. MQGMO_LOGICAL_ORDER cannot be used for the first retrieve, but can thereafter.

Restart processing looks like this:

```

MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
if (Reason == MQRC_NO_MSG_AVAILABLE)
  /* Proceed to normal processing */
  ...

else
  /* Group was terminated prematurely */
  /* The next message on the group must be retrieved by matching
  the sequence number and group ID with those retrieved from the
  status information. */
  GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT | MQGMO_WAIT
  MQGET GMO.MatchOptions = MQMO_MATCH_GROUP_ID | MQMO_MATCH_MSG_SEQ_NUMBER,
    MQMD.GroupId      = value from Status message,
    MQMD.MsgSeqNumber = value from Status message plus 1
  msgs = 1
  /* Process this message */
  ...

  /* Now normal processing is resumed */
  /* Retrieve remaining messages in the group */
  do while ( GroupStatus == MQGS_MSG_IN_GROUP )

    /* Process up to 4 messages in the group */
    GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT | MQGMO_WAIT
                  | MQGMO_LOGICAL_ORDER
    do while ( (GroupStatus == MQGS_MSG_IN_GROUP) && (msgs < 4) )
      MQGET
      msgs = msgs + 1
      /* Process this message */
      ...

    /* Have retrieved last message or 4 messages */
    /* Update status message if not last in group */
    MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
    if ( GroupStatus == MQGS_MSG_IN_GROUP )
      StatusInfo = GroupId,MsgSeqNumber from MQMD
      MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
    MQCMIT
    msgs = 0

```

Getting a particular message:

There are a number of ways of getting a particular message from a queue. These are: selecting on the *MsgId* and *CorrelId*, selecting on the *GroupId*, *MsgSeqNumber* and *Offset*, and selecting on the *MsgToken*. You can also use a selection string when you open the queue.

To get a particular message from a queue, use the *MsgId* and *CorrelId* fields of the MQMD structure. However, applications can explicitly set these fields, so the values that you specify might not identify a unique message. Table 20 shows which message is retrieved for the possible settings of these fields. These fields are ignored on input if you specify MQGMO_MSG_UNDER_CURSOR in the **GetMsgOpts** parameter of the MQGET call.

Table 20. Using message and correlation identifiers

To retrieve ...	<i>MsgId</i>	<i>CorrelId</i>
First message in the queue	MQMI_NONE	MQCI_NONE
First message that matches <i>MsgId</i>	Nonzero	MQCI_NONE
First message that matches <i>CorrelId</i>	MQMI_NONE	Nonzero
First message that matches both <i>MsgId</i> and <i>CorrelId</i>	Nonzero	Nonzero

In each case, *first* means the first message that satisfies the selection criteria (unless MQGMO_BROWSE_NEXT is specified, when it means the *next* message in the sequence satisfying the selection criteria).

On return, the MQGET call sets the *MsgId* and *CorrelId* fields to the message and correlation identifiers of the message returned, if any.

If you set the *Version* field of the MQMD structure to 2, you can use the *GroupId*, *MsgSeqNumber*, and *Offset* fields. Table 21 shows which message is retrieved for the possible settings of these fields.

Table 21. Using the group identifier

To retrieve ...	Match options
First message in the queue	MQMO_NONE
First message that matches <i>MsgId</i>	MQMO_MATCH_MSG_ID
First message that matches <i>CorrelId</i>	MQMO_MATCH_CORREL_ID
First message that matches <i>GroupId</i>	MQMO_MATCH_GROUP_ID
First message that matches <i>MsgSeqNumber</i>	MQMO_MATCH_MSG_SEQ_NUMBER
First message that matches <i>MsgToken</i>	MQMO_MATCH_MSG_TOKEN
First message that matches <i>Offset</i>	MQMO_MATCH_OFFSET

Notes:

1. MQMO_MATCH_XXX implies that the XXX field in the MQMD structure is set to the value to be matched.
2. The MQMO flags can be used in combination. For example, MQMO_MATCH_GROUP_ID, MQMO_MATCH_MSG_SEQ_NUMBER, and MQMO_MATCH_OFFSET can be used together to give the segment identified by the *GroupId*, *MsgSeqNumber*, and *Offset* fields.
3. If you specify MQGMO_LOGICAL_ORDER, the message that you are trying to retrieve is affected because the option depends on state information controlled for the queue handle. For information about this, see "Logical and physical ordering" on page 134 and Options.

The MQGET call usually retrieves the first message from a queue. If you specify a particular message when you use the MQGET call, the queue manager must search the queue until it finds that message. This can affect the performance of your application.

If you are using Version 2 or later of the MQGMO structure and do not specify the MQMO_MATCH_MSG_ID or MQMO_MATCH_CORREL_ID flags, you do not need to reset the *MsgId* or *CorrelId* fields between MQGETs.

z/OS On IBM MQ for z/OS, the queue attribute *IndexType* can be used to increase the speed of MQGET operations on the queue. For more information, see “Type of index” on page 152.

You can get a specific message from a queue by specifying its *MsgToken* and the *MatchOption* MQMO_MATCH_MSG_TOKEN in the MQGMO structure. The *MsgToken* is returned by the MQPUT call that originally put that message on the queue, or by previous MQGET operations and remains constant unless the queue manager is restarted.

If you are interested in only a subset of messages on the queue, you can specify which messages you want to process by using a selection string with the MQOPEN or MQSUB call. MQGET then retrieves the next message that satisfies that selection string. For more information about selection strings, see “Selectors” on page 22.

Improving performance of non-persistent messages:

When a client requires a message from a server, it sends a request to the server. It sends a separate request for each of the messages it consumes. To improve the performance of a client consuming non-persistent messages by avoiding having to send these request messages, a client can be configured to use *read ahead*. Read ahead allows messages to be sent to a client without an application having to request them.

When read ahead is enabled, messages are sent to a memory buffer on the client called the *read ahead buffer*. The client will have a read ahead buffer for each queue it has open with read ahead enabled. The messages in the read ahead buffer are not persisted. The client periodically updates the server with information about the amount of data it has consumed.

When you call MQOPEN with MQOO_READ_AHEAD, the IBM MQ client only enables read-ahead if certain conditions are met. These conditions include:

- Both the client and remote queue manager must be at WebSphere MQ Version 7 or later.
- The client application must be compiled and linked against the threaded IBM MQ MQI client libraries.
- The client channel must be using TCP/IP protocol
- The channel must have a non-zero *SharingConversations* (SHARECNV) setting in both the client and server channel definitions.

Using read ahead can improve performance when consuming non-persistent messages from a client application. This performance improvement is available to both MQI and JMS applications. Client applications using MQGET or asynchronous consumption will benefit from the performance improvements when consuming non-persistent messages.

Not all client application designs are suited to using read ahead as not all options are supported for use with read ahead and some options are required to be consistent between MQGET calls when read ahead is enabled. If a client alters its selection criteria between MQGET calls, messages being stored in the read ahead buffer will remain stranded in the client read ahead buffer.

If a backlog of stranded messages with the previous selection criteria are no longer required, a configurable purge interval can be set on the client to automatically purge these messages from the client. The purge interval is one of a group of read ahead tuning options determined by the client. It is possible to tune these options to meet your requirements.

If a client application is restarted, messages in the read ahead buffer can be lost. Conversely, a message that has been moved to a read ahead buffer could then be deleted from the underlying queue; this does not result in it being removed from the buffer, so an MQGET call using read ahead can return a message that no longer exists.

Read ahead is only performed for client bindings. The attribute is ignored for all other bindings.

Read ahead has no effect on triggering. No trigger message is generated when a message is read ahead by the client. Read ahead does not generate accounting and statistics information when it is enabled.

Using read ahead with publish subscribe messaging

When a subscribing application specifies a destination queue to which publications are sent, the DEFREADA value of the specified queue is used as the default read ahead value.

When a subscribing application requests that IBM MQ manages the destination to which publications are sent, a managed queue is created as a dynamic queue based upon a predefined model queue. It is the DEFREADA value of the model queue that is used as the default read ahead value. The default model queues SYSTEM.DURABLE.PUBLICATIONS.MODEL or SYSTEM.NONDURABLE.PUBLICATIONS.MODEL are used unless a model queue is defined for this or a parent topic.

Related concepts:

“Tuning performance for nonpersistent messages on AIX” on page 151

If you are using AIX V5.3 or later, consider setting your tuning parameter to use full performance for nonpersistent messages.

Related tasks:

“Enabling and disabling read ahead” on page 150

By default read ahead is disabled. You can enable read ahead at queue or application level.

Related reference:

“MQGET options and read ahead”

Not all MQGET options are supported when read ahead is enabled; some options are required to be consistent between MQGET calls.

MQGET options and read ahead:

Not all MQGET options are supported when read ahead is enabled; some options are required to be consistent between MQGET calls.

When you call MQOPEN with MQOO_READ_AHEAD, the IBM MQ client only enables read-ahead if certain conditions are met. These conditions include:

- Both the client and remote queue manager must be at WebSphere MQ Version 7 or later.
- The client application must be compiled and linked against the threaded IBM MQ MQI client libraries.
- The client channel must be using TCP/IP protocol
- The channel must have a non-zero SharingConversations (SHARECNV) setting in both the client and server channel definitions.

The following table indicates which options are supported for use with read ahead and whether they can be altered between MQGET calls.

Table 22. MQGET options and read ahead

	Permitted when read ahead is enabled and can be altered between MQGET calls ⁵	Permitted when read ahead is enabled but cannot be altered between MQGET calls ¹	MQGET Options that are not permitted when read ahead is enabled ²
MQGET MQMD values	MsgId ³ CorrelId ³	Encoding CodedCharSetId	
MQGET MQGMO Options	<ul style="list-style-type: none"> • MQGMO_NO_WAIT • MQGMO_BROWSE_MESSAGE_UNDER_CURSOR • MQGMO_BROWSE_FIRST • MQGMO_BROWSE_NEXT • MQGMO_FAIL_IF_QUIESCING 	<ul style="list-style-type: none"> • MQGMO_SYNCPOINT_IF_PERSISTENT • MQGMO_NO_SYNCPOINT • MQGMO_ACCEPT_TRUNCATED_MSG • MQGMO_CONVERT 	<ul style="list-style-type: none"> • MQGMO_SET_SIGNAL • MQGMO_SYNCPOINT • MQGMO_MARK_SKIP_BACKOUT • MQGMO_MSG_UNDER_CURSOR ⁴ • MQGMO_LOCK • MQGMO_UNLOCK • MQGMO_LOGICAL_ORDER • MQGMO_COMPLETE_MSG • MQGMO_ALL_MSGS_AVAILABLE • MQGMO_ALL_SEGMENTS_AVAILABLE

Notes:

1. If these options are altered between MQGET calls, an MQRC_OPTIONS_CHANGED reason code is returned.
2. If these options are specified on the first MQGET call then read ahead is disabled. If these options are specified on a subsequent MQGET call a reason code MQRC_OPTIONS_ERROR is returned.
3. If a client application alters MsgId and CorrelId values between MQGET calls, messages with the previous values might already have been sent to the client and will remain in the client read ahead buffer until consumed (or automatically purged).
4. MQGMO_MSG_UNDER_CURSOR is not possible with read ahead. Read ahead is disabled when both MQOO_BROWSE and one of the MQOO_INPUT_SHARED or MQOO_INPUT_EXCLUSIVE options are specified when opening the queue.
5. When read ahead is enabled, the first MQGET determines whether messages are to be browsed or got from a queue. If the client application then uses MQGET with changed options, such as attempting to browse following an initial get, or attempting to get following an initial browse, an MQRC_OPTIONS_CHANGED reason code is returned.

If a client alters its selection criteria between MQGET calls, messages being stored in the read ahead buffer that match the initial selection criteria are not consumed by the client application, and remain stranded in the client read ahead buffer. In situations where the client read ahead buffer contains many stranded messages, the benefits associated with read ahead are lost and a separate request to the server is required for each message consumed. To determine whether read ahead is being used efficiently you can use the connection status parameter, READA.

Read ahead can be inhibited when requested by an application due to incompatible options specified on the first MQGET call. In this situation the connection status shows read ahead as being inhibited.

If, because of these restrictions on MQGET, you decide that a client application design is not suited to read ahead, specify the MQOPEN option MQOO_READ_AHEAD_NO. Alternatively set the default read ahead value of the queue being opened altered to either NO or DISABLED.

Enabling and disabling read ahead:

By default read ahead is disabled. You can enable read ahead at queue or application level.

About this task

When you call MQOPEN with MQOO_READ_AHEAD, the IBM MQ client only enables read-ahead if certain conditions are met. These conditions include:

- Both the client and remote queue manager must be at WebSphere MQ Version 7 or later.
- The client application must be compiled and linked against the threaded IBM MQ MQI client libraries.
- The client channel must be using TCP/IP protocol
- The channel must have a non-zero SharingConversations (SHARECNV) setting in both the client and server channel definitions.

To enable read ahead:

- To configure read ahead at the queue level set the queue attribute, DEFREADA to YES.
- To configure read ahead at the application level:
 - to use read ahead wherever possible use the MQOO_READ_AHEAD option on the MQOPEN function call. It is not possible for the client application to use read ahead if the DEFREADA queue attribute has been set to DISABLED.
 - to use read ahead only when read ahead is enabled on a queue, use the MQOO_READ_AHEAD_AS_Q_DEF option on the MQOPEN function call.

If a client application design is not suited to read ahead you can disable it:

- at the queue level by setting the queue attribute, DEFREADA to NO if you do not want read ahead to be used unless it is requested by a client application, or DISABLED if you do not want read ahead to be used regardless of whether read ahead is required by a client application.
- at the application level by using the MQOO_NO_READ_AHEAD option on the MQOPEN function call.

Two MQCLOSE options allow you to configure what happens to any messages that are being stored in the read ahead buffer if the queue is closed.

- Use MQCO_IMMEDIATE to discard messages in the read ahead buffer.
- Use MQCO QUIESCE to ensure that messages in the read ahead buffer are consumed by the application before the queue is closed. When MQCLOSE with the MQCO QUIESCE is issued and there are messages remaining on the read ahead buffer, MQRC_READ_AHEAD_MSGS returns with MQCC_WARNING.

Tuning performance for nonpersistent messages on AIX:

If you are using AIX V5.3 or later, consider setting your tuning parameter to use full performance for nonpersistent messages.

To set the tuning parameter so that it takes effect immediately, issue the following command as a root user:

```
/usr/sbin/iioo -o j2_nPagesPerWriteBehindCluster=0
```

To set the tuning parameter so that it takes effect immediately and persists over reboots, issue the following command as a root user:

```
/usr/sbin/iioo -p -o j2_nPagesPerWriteBehindCluster=0
```

Normally, nonpersistent messages are kept only in memory, but there are circumstances where AIX can schedule nonpersistent messages to be written to disk. Messages scheduled to be written to disk are unavailable for MQGET until the disk write completes. The suggested tuning command varies this threshold; instead of scheduling messages to be written to disk when 16 kilobytes of data are queued, the write-to-disk occurs only when real storage on the machine becomes close to full. This is a global alteration and might affect other software components.

On AIX, when using multithreaded applications and especially when running on machines with multiple processors, we strongly recommend setting AIXTHREAD_SCOPE=S in the mqm ID .profile or setting AIXTHREAD_SCOPE=S in the environment before starting the application, for better performance and more solid scheduling. For example:

```
export AIXTHREAD_SCOPE=S
```

Setting AIXTHREAD_SCOPE=S means that user threads created with default attributes are placed into system-wide contention scope. If a user thread is created with system-wide contention scope, it is bound to a kernel thread and it is scheduled by the kernel. The underlying kernel thread is not shared with any other user thread.

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, you can increase the process limit for the number of file descriptors. To do so, alter the nfiles attribute in /etc/security/limits to 10,000 for the mqm user ID or in the default stanza.

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
```

Type of index:

The queue attribute, *IndexType*, specifies the type of index that the queue manager maintains to increase the speed of MQGET operations on the queue.

Note: Supported only on IBM MQ for z/OS.

You have five options:

Value	Description
NONE	No index is maintained. Use this when retrieving messages sequentially (see “Priority” on page 134).
GROUPID	An index of group identifiers is maintained. You must use this index type if you want logical ordering of message groups (see “Logical and physical ordering” on page 134).
MSGID	An index of message identifiers is maintained. Use this when retrieving messages using the <i>MsgId</i> field as a selection criterion on the MQGET call (see “Getting a particular message” on page 146).
MSGTOKEN	An index of message tokens is maintained.
CORRELID	An index of correlation identifiers is maintained. Use this when retrieving messages using the <i>CorrelId</i> field as a selection criterion on the MQGET call (see “Getting a particular message” on page 146).

Note:

1. If you are indexing using the MSGID option or CORRELID option, set the relative **MsgId** or **CorrelId** parameters in the MQMD. It is not beneficial to set both.
2. Browse uses the index mechanism to find a message if a queue matches all the following conditions:
 - It has index type MSGID, CORRELID, or GROUPID
 - It is browsed with the same type of id
 - It has messages of only one priority
3. Avoid queues (indexed by *MsgId* or *CorrelId*) containing thousands of messages because this affects restart time. (This does not apply to nonpersistent messages as they are deleted at restart.)
4. MSGTOKEN is used to define queues managed by the z/OS workload manager.

For a full description of the **IndexType** attribute, see *IndexType*. For further information on the **IndexType** attribute, see “Design and performance considerations for z/OS applications” on page 53.

Handling messages greater than 4 MB long:

Messages can be too large for the application, queue, or queue manager. Depending on the environment, IBM MQ provides a number of ways of dealing with messages that are longer than 4 MB.

You can increase the **MaxMsgLength** attribute up to 100 MB on all IBM MQ systems at V6 or later. Set this value to reflect the size of the messages using the queue. On IBM MQ systems other than IBM MQ for z/OS, you can also:

1. Use segmented messages. (Messages can be segmented by either the application or the queue manager.)
2. Use reference messages.

Each of these approaches is described in the remainder of this section.

Increasing the maximum message length

The **MaxMsgLength** queue manager attribute defines the maximum length of a message that can be handled by a queue manager. Similarly, the **MaxMsgLength** queue attribute is the maximum length of a message that can be handled by a queue. The default maximum message length supported depends on the environment in which you are working.

If you are handling large messages, you can alter these attributes independently. You can set the queue manager attribute value in the range 32768 bytes through 100 MB; you can set the queue attribute value in the range 0 through 100 MB.


After changing one or both of the **MaxMsgLength** attributes, restart your applications and channels to ensure that the changes take effect.

When these changes are made, the message length must be less than or equal to both the queue and the queue manager **MaxMsgLength** attributes. However, existing messages might be longer than either attribute.

If the message is too big for the queue, `MQRC_MSG_TOO_BIG_FOR_Q` is returned. Similarly, if the message is too big for the queue manager, `MQRC_MSG_TOO_BIG_FOR_Q_MGR` is returned.

This method of handling large messages is easy and convenient. However, consider the following factors before using it:

- Uniformity among queue managers is reduced. The maximum size of message data is determined by the *MaxMsgLength* for each queue (including transmission queues) on which the message will be put. This value is often defaulted to the queue manager's *MaxMsgLength*, especially for transmission queues. This makes it difficult to predict whether a message is too large when it is to travel to a remote queue manager.
- Usage of system resources is increased. For example, applications need larger buffers, and on some platforms, there might be increased usage of shared storage. Queue storage should be affected only if actually required for larger messages.
- Channel batching is affected. A large message still counts as just one message towards the batch count but needs longer to transmit, thereby increasing response times for other messages.

Message segmentation: 

Use this information to learn about segmenting messages. This feature is not supported on IBM MQ for z/OS or by applications using IBM MQ classes for JMS.

Increasing the maximum message length as explained in topic “Increasing the maximum message length” has some negative implications. Also, it can still result in the message being too large for the queue or queue manager. In these cases, you can segment a message. For information about segments, see “Message groups” on page 36.


The next sections look at common uses for segmenting messages. For putting and destructively getting, it is assumed that the `MQPUT` or `MQGET` calls *always* operate within a unit of work. Always consider using this technique to reduce the possibility of incomplete groups being present in the network. Single-phase commit by the queue manager is assumed, but other coordination techniques are equally valid.

Also, in the getting applications, it is assumed that if multiple servers are processing the same queue, each server executes similar code, so that one server never fails to find a message or segment that it expects to be there (because it had specified `MQGMO_ALL_MSGS_AVAILABLE` or `MQGMO_ALL_SEGMENTS_AVAILABLE` earlier).

Putting and getting a segmented message that spans units of work

You can put and get a segmented message that spans a unit of work in a similar way to “Putting and getting a group that spans units of work” on page 143.

You cannot, however, put or get segmented messages in a global unit of work.

Segmentation and reassembly by queue manager: 

This is the simplest scenario, in which one application puts a message to be retrieved by another. The message might be large: not too large for either the putting or the getting application to handle in a single buffer, but too large for the queue manager or a queue on which the message is to be put.

The only changes necessary for these applications are for the putting application to authorize the queue manager to perform segmentation if necessary:

```
PMO.Options = (existing options)
MD.MsgFlags = MQMF_SEGMENTATION_ALLOWED
MD.Version = MQMD_VERSION_2
memcpy(MD.GroupId, MQGI_NONE, MQ_GROUP_ID_LENGTH)
MQPUT
```

and for the getting application to ask the queue manager to reassemble the message if it has been segmented:

```
GMO.Options = MQGMO_COMPLETE_MSG | (existing options)
MQGET
```

In this simplest scenario, the application must reset the GroupId field to MQGI_NONE before the MQPUT call, so that the queue manager can generate a unique group identifier for each message. If this is not done, unrelated messages can have the same group identifier, which might subsequently lead to incorrect processing.

The application buffer must be large enough to contain the reassembled message (unless you include the MQGMO_ACCEPT_TRUNCATED_MSG option).

If the MAXMSGLEN attribute of a queue is to be modified to accommodate message segmentation, then consider:

- The minimum message segment supported on a local queue is 16 bytes.
- For a transmission queue, MAXMSGLEN must also include the space required for headers. Consider using a value at least 4000 bytes larger than the maximum expected length of user data in any message segment that could be put on a transmission queue.

If data conversion is necessary, the getting application might have to do it by specifying MQGMO_CONVERT. This should be straightforward because the data conversion exit is presented with the complete message. Do not attempt to convert data in a sender channel if the message is segmented, and the format of the data is such that the data-conversion exit cannot carry out the conversion on incomplete data.

Application segmentation: 

Application segmentation is used when queue manager segmentation is not adequate, or when applications require data conversion with specific segment boundaries.

Application segmentation is used for two main reasons:

1. Queue manager segmentation alone is not adequate because the message is too large to be handled in a single buffer by the applications.
2. Data conversion must be performed by sender channels, and the format is such that the putting application must stipulate where the segment boundaries are to be in order for conversion of an individual segment to be possible.

However, if data conversion is not an issue, or if the getting application always uses `MQGMO_COMPLETE_MSG`, queue manager segmentation can also be allowed by specifying `MQMF_SEGMENTATION_ALLOWED`. In our example, the application segments the message into four segments:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

MQPUT MD.MsgFlags = MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_LAST_SEGMENT

MQCMIT
```

If you do not use `MQPMO_LOGICAL_ORDER`, the application must set the *Offset* and the length of each segment. In this case, logical state is not maintained automatically.

The getting application cannot guarantee to have a buffer large enough to hold any reassembled message. It must therefore be prepared to process segments individually.

For messages that are segmented, this application does not want to start processing one segment until all the segments that constitute the logical message are present. `MQGMO_ALL_SEGMENTS_AVAILABLE` is therefore specified for the first segment. If you specify `MQGMO_LOGICAL_ORDER` and there is a current logical message, `MQGMO_ALL_SEGMENTS_AVAILABLE` is ignored.

After the first segment of a logical message has been retrieved, use `MQGMO_LOGICAL_ORDER` to ensure that the remaining segments of the logical message are retrieved in order.

No consideration is given to messages within different groups. If such messages occur, they are processed in the order in which the first segment of each message occurs on the queue.

```
GMO.Options = MQGMO_SYNCPOINT | MQGMO_LOGICAL_ORDER
              | MQGMO_ALL_SEGMENTS_AVAILABLE | MQGMO_WAIT
do while ( SegmentStatus == MQSS_SEGMENT )
  MQGET
  /* Process each remaining segment of the logical message */
  ...

MQCMIT
```

Application segmentation of logical messages:

Multi

The messages must be maintained in logical order in a group, and some or all of them might be so large that they require application segmentation.

In our example, a group of four logical messages is to be put. All but the third message are large, and require segmentation, which is performed by the putting application:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_LAST_SEGMENT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_LAST_SEGMENT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP

MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_LAST_SEGMENT

MQCMIT
```

In the getting application, MQGMO_ALL_MSGS_AVAILABLE is specified on the first MQGET. This means that no messages or segments of a group are retrieved until the entire group is available. When the first physical message of a group has been retrieved, MQGMO_LOGICAL_ORDER is used to ensure that the segments and messages of the group are retrieved in order:

```
GMO.Options = MQGMO_SYNCPOINT | MQGMO_LOGICAL_ORDER
              | MQGMO_ALL_MSGS_AVAILABLE | MQGMO_WAIT

do while ( (GroupStatus  != MQGS_LAST_MSG_IN_GROUP) ||
           (SegmentStatus != MQGS_LAST_SEGMENT) )
  MQGET
  /* Process a segment or complete logical message. Use the GroupStatus
   * and SegmentStatus information to see what has been returned */
  ...
MQCMIT
```

Note: If you specify MQGMO_LOGICAL_ORDER and there is a current group, MQGMO_ALL_MSGS_AVAILABLE is ignored.

Reference messages:

Use this information to learn more about reference messages.

Note: Not supported in IBM MQ for z/OS.

This method allows a large object to be transferred from one node to another without storing the object on IBM MQ queues at either the source or the destination nodes. This is of particular benefit when the data exists in another form, for example, for mail applications.

To do this, you specify a message exit at both ends of a channel. For information about how to do this, see “Channel message exit programs” on page 360.

IBM MQ defines the format of a reference message header (MQRMH). See MQRMH for a description of this. This is recognized with a defined format name and might be followed by actual data.

To initiate transfer of a large object, an application can put a message consisting of a reference message header with no data following it. As this message leaves the node, the message exit retrieves the object in an appropriate way and appends it to the reference message. It then returns the message (now larger than before) to the sending Message Channel Agent for transmission to the receiving MCA.

Another message exit is configured at the receiving MCA. When this message exit receives one of these messages, it creates the object using the object data that was appended and passes on the reference message *without* it. The reference message can now be received by an application and this application knows that the object (or at least the portion of it represented by this reference message) has been created at this node.

The maximum amount of object data that a sending message exit can append to the reference message is limited by the negotiated maximum message length for the channel. The exit can return only a single message to the MCA for each message that it is passed, so the putting application can put several messages to cause one object to be transferred. Each message must identify the *logical* length and offset of the object that is to be appended to it. However, in cases where it is not possible to know the total size of the object or the maximum size allowed by the channel, design the sending message exit so that the putting application just puts a single message, and the exit itself puts the next message on the transmission queue when it has appended as much data as it can to the message it has been passed.

Before using this method of dealing with large messages, consider the following points:

- The MCA and the message exit run under an IBM MQ user ID. The message exit (and therefore, the user ID) needs to access the object to either retrieve it at the sending end or create it at the receiving end; this might only be feasible in cases where the object is universally accessible. This raises a security issue.
- If the reference message with bulk data appended to it must travel through several queue managers before reaching its destination, the bulk data is present on IBM MQ queues at the intervening nodes. However, no special support or exits need to be provided in these cases.
- Designing your message exit is made difficult if rerouting or dead-letter queuing is allowed. In these cases, the portions of the object might arrive out of order.
- When a reference message arrives at its destination, the receiving message exit creates the object. However, this is not synchronized with the MCA's unit of work, so if the batch is backed out, another reference message containing this same portion of the object will arrive in a later batch, and the message exit might attempt to re-create the same portion of the object. If the object is, for example, a series of database updates, this might be unacceptable. If so, the message exit must keep a log of which updates have been applied; this might require the use of an IBM MQ queue.

- Depending on the characteristics of the object type, the message exits and applications might need to cooperate in maintaining use counts, so that the object can be deleted when it is no longer needed. An instance identifier might also be required; a field is provided for this in the reference message header (see MQRMH).
- If a reference message is put as a distribution list, the object must be retrievable for each resulting distribution list or individual destination at that node. You might need to maintain use counts. Also consider the possibility that a node might be the final node for some of the destinations in the list, but an intermediate node for others.
- Bulk data is not typically converted. This is because conversion takes place *before* the message exit is invoked. For this reason, conversion must not be requested on the originating sender channel. If the reference message passes through an intermediate node, the bulk data is converted when sent from the intermediate node, if requested.
- Reference messages cannot be segmented.

Using the MQRMH and MQMD structures

See MQRMH and MQMD for a description of the fields in the reference message header and the message descriptor.

In the MQMD structure, set the *Format* field to MQFMT_REF_MSG_HEADER. The MQHREF format, when requested on MQGET, is converted automatically by IBM MQ along with any bulk data that follows.

Here is an example of the use of the *DataLogicalOffset* and *DataLogicalLength* fields of the MQRMH:

A putting application might put a reference message with:

- No physical data
- *DataLogicalLength* = 0 (this message represents the entire object)
- *DataLogicalOffset* = 0.

Assuming that the object is 70 000 bytes long, the sending message exit sends the first 40 000 bytes along the channel in a reference message containing:

- 40 000 bytes of physical data following the MQRMH
- *DataLogicalLength* = 40000
- *DataLogicalOffset* = 0 (from the start of the object).

It then places another message on the transmission queue containing:

- No physical data
- *DataLogicalLength* = 0 (to the end of the object). You could specify a value of 30 000 here.
- *DataLogicalOffset* = 40000 (starting from this point).

When this message exit is seen by the sending message exit, the remaining 30,000 bytes of data are appended, and the fields are set to:

- 30,000 bytes of physical data following the MQRMH
- *DataLogicalLength* = 30000
- *DataLogicalOffset* = 40000 (starting from this point).

The MQRMHF_LAST flag is also set.

For a description of the sample programs provided for the use of reference messages, see “Using the sample programs on Multiplatforms” on page 457.

Waiting for messages:

If you want a program to wait until a message arrives on a queue, specify the `MQGMO_WAIT` option in the *Options* field of the `MQGMO` structure.


Use the *WaitInterval* field of the `MQGMO` structure to specify the maximum time (in milliseconds) that you want an `MQGET` call to wait for a message to arrive on a queue.


If the message does not arrive within this time, the `MQGET` call completes with the `MQRC_NO_MSG_AVAILABLE` reason code.

You can specify an unlimited wait interval using the constant `MQWI_UNLIMITED` in the *WaitInterval* field. However, events outside your control could cause your program to wait for a long time, so use this constant with caution. IMS applications must not specify an unlimited wait interval because this would prevent the IMS system terminating. (When IMS terminates, it requires all dependent regions to end.) Instead, IMS applications can specify a finite wait interval; then, if the call completes without retrieving a message after that interval, issue another `MQGET` call with the wait option.

Note: If more than one program is waiting on the same shared queue to *remove* a message, only one program is activated by a message arriving. However, if more than one program is waiting to browse a message, all the programs can be activated. For more information, see the description of the *Options* field of the `MQGMO` structure in `MQGMO`.

If the state of the queue or the queue manager changes before the wait interval expires, the following actions occur:

- If the queue manager enters the quiescing state, and you used the `MQGMO_FAIL_IF QUIESCING` option, the wait is canceled and the `MQGET` call completes with the `MQRC_Q_MGR QUIESCING` reason code. Without this option, the call remains waiting.
-  On z/OS, if the connection (for a CICS or IMS application) enters the quiescing state, and you used the `MQGMO_FAIL_IF QUIESCING` option, the wait is canceled and the `MQGET` call completes with the `MQRC_CONN QUIESCING` reason code. Without this option, the call remains waiting.
- If the queue manager is forced to stop, or is canceled, the `MQGET` call completes with either the `MQRC_Q_MGR STOPPING` or the `MQRC_CONNECTION_BROKEN` reason code.
- If the attributes of the queue (or a queue to which the queue name resolves) are changed so that get requests are now inhibited, the wait is canceled and the `MQGET` call completes with the `MQRC_GET_INHIBITED` reason code.
- If the attributes of the queue (or a queue to which the queue name resolves) are changed in such a way that the `FORCE` option is required, the wait is canceled and the `MQGET` call completes with the `MQRC_OBJECT_CHANGED` reason code.

 If you want your application to wait on more than one queue, use the signal facility of IBM MQ for z/OS (see “Signaling” on page 160). For more information about the circumstances in which these actions occur, see `MQGMO`.

Signaling:

Signaling is supported only on IBM MQ for z/OS.

Signaling is an option on the MQGET call to allow the operating system to notify (or *signal*) a program when an expected message arrives on a queue. This is like the *get with wait* function described in topic “Waiting for messages” on page 159 because it allows your program to continue with other work while waiting for the signal. However, if you use signaling, you can free the application thread and rely on the operating system to notify the program when a message arrives.

To set a signal

To set a signal, do the following in the MQGMO structure that you use on your MQGET call:

1. Set the MQGMO_SET_SIGNAL option in the *Options* field.
2. Set the maximum life of the signal in the *WaitInterval* field. This sets the length of time (in milliseconds) for which you want IBM MQ to monitor the queue. Use the MQWI_UNLIMITED value to specify an unlimited life.

Note: IMS applications must not specify an unlimited wait interval because this would prevent the IMS system from terminating. (When IMS terminates, it requires all dependent regions to end.) Instead, IMS applications can examine the state of the ECB at regular intervals (see step 3). A program can have signals set on several queue handles at the same time:

3. Specify the address of the *Event Control Block* (ECB) in the *Signal1* field. This notifies you of the result of your signal. The ECB storage must remain available until the queue is closed.

Note: You cannot use the MQGMO_SET_SIGNAL option with the MQGMO_WAIT option.

When the message arrives

When a suitable message arrives, a completion code is returned to the ECB.

The completion code describes one of the following:

- The message that you set the signal for has arrived on the queue. The message is not reserved for the program that requested a signal, so the program must issue an MQGET call again to get the message.

Note: Another application could get the message in the time between your receiving the signal and issuing another MQGET call.

- The wait interval you set has expired and the message you set the signal for did not arrive on the queue. IBM MQ has canceled the signal.
- The signal has been canceled. This happens, for example, if the queue manager stops, or the attribute of the queue is changed, so that MQGET calls are no longer allowed.

When a suitable message is already on the queue, the MQGET call completes in the same way as an MQGET call without signaling. Also, if an error is detected immediately, the call completes and the return codes are set.

When the call is accepted and no message is immediately available, control is returned to the program so that it can continue with other work. None of the output fields in the message descriptor are set, but the **CompCode** parameter is set to MQCC_WARNING and the **Reason** parameter is set to MQRC_SIGNAL_REQUEST_ACCEPTED.

For information about what IBM MQ can return to your application when it makes an MQGET call using signaling, see MQGET.

If the program has no other work to do while it is waiting for the ECB to be posted, it can wait for the ECB using:

- For a CICS Transaction Server for z/OS program, the EXEC CICS WAIT EXTERNAL command
- For batch and IMS programs, the z/OS WAIT macro

If the state of the queue or the queue manager changes while the signal is set (that is, the ECB has not yet been posted), the following actions occur:

- If the queue manager enters the quiescing state, and you used the MQGMO_FAIL_IF_QUIESCING option, the signal is canceled. The ECB is posted with the MQEC_Q_MGR_QUIESCING completion code. Without this option, the signal remains set.
- If the queue manager is forced to stop, or is canceled, the signal is canceled. The signal is delivered with the MQEC_WAIT_CANCELED completion code.
- If the attributes of the queue (or a queue to which the queue name resolves) are changed so that get requests are now inhibited, the signal is canceled. The signal is delivered with the MQEC_WAIT_CANCELED completion code.

Note:

1. If more than one program has set a signal on the same shared queue to remove a message, only one program is activated by a message arriving. However, if more than one program is waiting to browse a message, all the programs can be activated. The rules that the queue manager follows when deciding which applications to activate are the same as those for waiting applications: for more information, see the description of the *Options* field of the MQGMO structure in MQGMO - Get-message options.
2. If there is more than one MQGET call waiting for the same message, with a mixture of wait and signal options, each waiting call is considered equally. For more information, see the description of the *Options* field of the MQGMO structure in MQGMO - Get-message options.
3. Under some conditions, it is possible both for an MQGET call to retrieve a message and for a signal (resulting from the arrival of the same message) to be delivered. This means that when your program issues another MQGET call (because the signal was delivered), there could be no message available. Design your program to test for this situation.

For information about how to set a signal, see the description of the MQGMO_SET_SIGNAL option and the *Signal1* field in Signal1.

Skipping backout:

You can prevent an application program from entering an *MQGET-error-backout* loop by specifying the **MQGMO_MARK_SKIP_BACKOUT** option on the MQGET call.

Note: Supported only on IBM MQ for z/OS.

As part of a unit of work, an application program can issue one or more MQGET calls to get messages from a queue. If the application program detects an error, it can back out the unit of work. This restores all the resources updated during that unit of work to the state that they were in before the unit of work started, and reinstates the messages retrieved by the MQGET calls.

Once reinstated, these messages are available to subsequent MQGET calls issued by the application program. In many cases, this does not cause a problem for the application program. However, in cases where the error leading to the backout cannot be circumvented, having the message reinstated on the queue can cause the application program to enter an *MQGET-error-backout* loop.

To avoid this problem, specify the MQGMO_MARK_SKIP_BACKOUT option on the MQGET call. This marks the MQGET request as not being involved in application-initiated backout; that is, it must not be

backed out. Use of this option means that when a backout occurs, updates to other resources are backed out as required, but the marked message is treated as if it had been retrieved under a new unit of work.

The application program must issue an IBM MQ call either to commit the new unit of work, or to back out the new unit of work. For example, the program can perform exception handling, such as informing the originator that the message has been discarded, and commit the unit of work so removing the message from the queue. If the new unit of work is backed out (for any reason) the message is reinstated on the queue.

Within a unit of work, there can be only one MQGET request marked as skipping backout; however, there can be several other messages that are not marked as skipping backout. Once a message has been marked as skipping backout, any further MQGET calls within the unit of work that specify MQGMO_MARK_SKIP_BACKOUT fail with reason code MQRC_SECOND_MARK_NOT_ALLOWED.

Note:

1. The marked message skips backout only if the unit of work containing it is terminated by an application request to back it out. If the unit of work is backed out for any other reason, the message is backed out onto the queue in the same way that it would be if it was not marked to skip backout.
2. Skip backout is not supported within Db2 stored procedures participating in units of work controlled by RRS. For example, an MQGET call with the MQGMO_MARK_SKIP_BACKOUT option will fail with the reason code MQRC_OPTION_ENVIRONMENT_ERROR.

Figure 16 on page 163 illustrates a typical sequence of steps that an application program might contain when an MQGET request is required to skip backout.

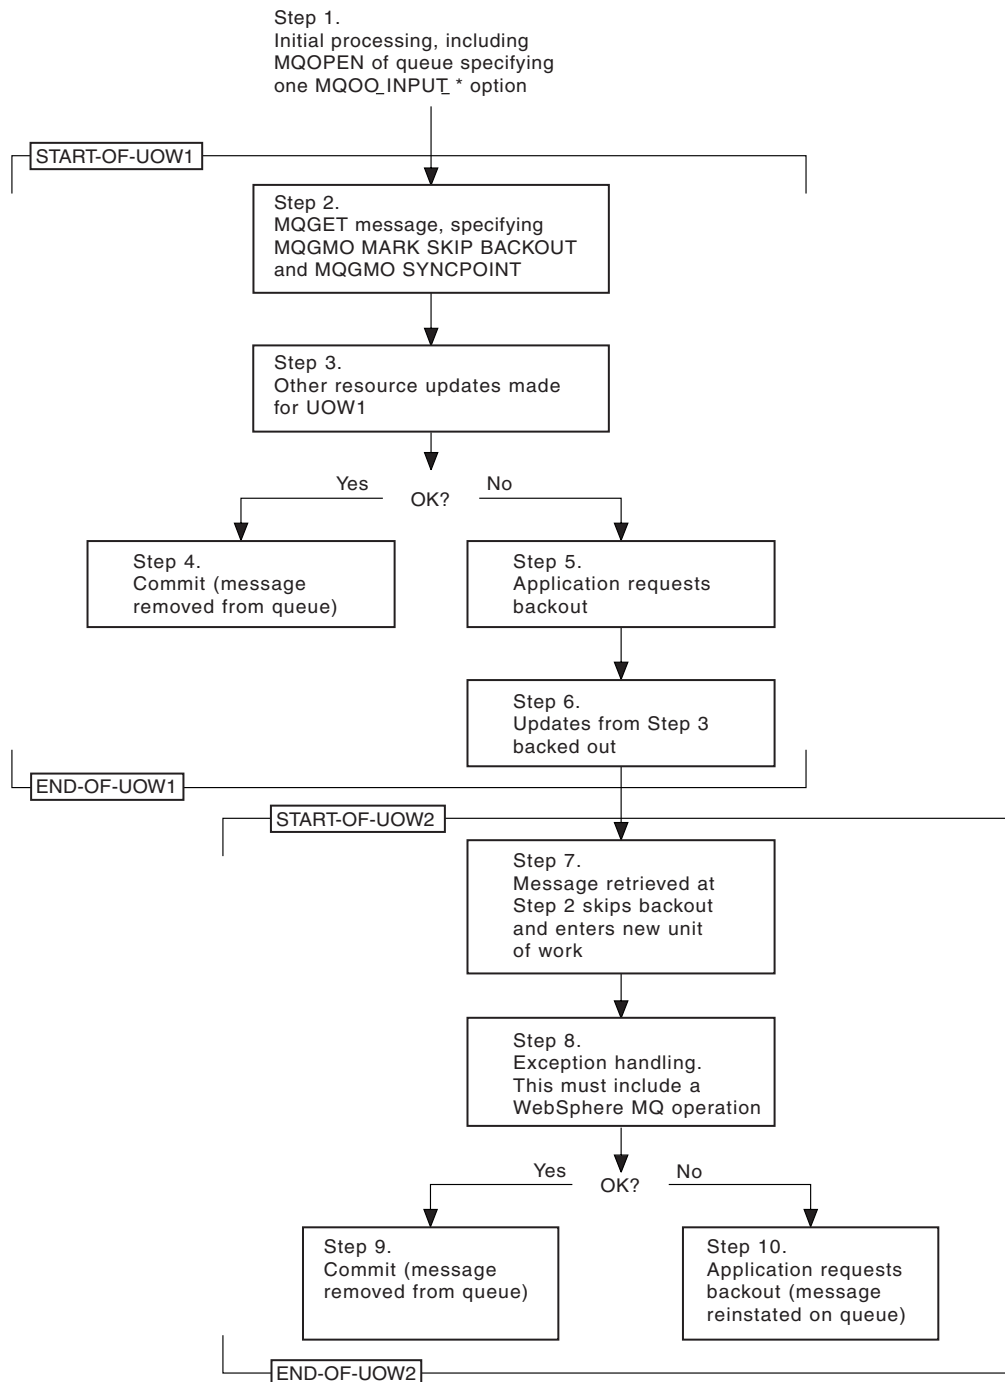


Figure 16. Skipping backout using MQGMO_MARK_SKIP_BACKOUT

The steps in Figure 16 are:

- Step 1** Initial processing occurs within the transaction, including an MQOPEN call to open the queue (specifying one of the MQOO_INPUT_* options in order to get messages from the queue in Step 2).
- Step 2** MQGET is called, with MQGMO_SYNCPOINT and MQGMO_MARK_SKIP_BACKOUT. MQGMO_SYNCPOINT is required because MQGET must be within a unit of work for MQGMO_MARK_SKIP_BACKOUT to be effective. In Figure 16 this unit of work is referred to as UOW1.

- Step 3** Other resource updates are made as part of UOW1. These can include further MQGET calls (issued without MQGMO_MARK_SKIP_BACKOUT).
- Step 4** All updates from Steps 2 and 3 complete as required. The application program commits the updates and UOW1 ends. The message retrieved in Step 2 is removed from the queue.
- Step 5** Some of the updates from Steps 2 and 3 do not complete as required. The application program requests that the updates made during these steps are backed out.
- Step 6** The updates made in Step 3 are backed out.
- Step 7** The MQGET request made in Step 2 skips backout and becomes part of a new unit of work, UOW2.
- Step 8** UOW2 performs exception handling in response to UOW1 being backed out. (For example, an MQPUT call to another queue, indicating that a problem occurred that caused UOW1 to be backed out.)
- Step 9** Step 8 completes as required, the application program commits the activity, and UOW2 ends. As the MQGET request is part of UOW2 (see Step 7), this commit causes the message to be removed from the queue.
- Step 10**
Step 8 does not complete as required and the application program backs out UOW2. Because the get message request is part of UOW2 (see Step 7), it too is backed out and reinstated on the queue. It is now available to further MQGET calls issued by this or another application program (in the same way as any other message on the queue).

Application data conversion:

When necessary, MCAs convert the message descriptor and header data into the required character set and encoding. Either end of the link (that is, the local MCA or the remote MCA) can do the conversion.

When an application puts messages on a queue, the local queue manager adds control information to the message descriptors to facilitate the control of the messages when they are processed by queue managers and MCAs. Depending on the environment, the message header data fields are created in the character set and encoding of the local system.

When you move messages between systems, you sometimes need to convert the application data into the character set and encoding required by the receiving system. This can be done either from within application programs on the receiving system or by the MCAs on the sending system. If data conversion is supported on the receiving system, use application programs to convert the application data, rather than depending on the conversion having already occurred at the sending system.

Application data is converted within an application program when you specify the MQGMO_CONVERT option in the *Options* field of the MQGMO structure passed to an MQGET call, and when *all* the following statements are true:

- The *CodedCharSetId* or *Encoding* fields set in the MQMD structure associated with the message on the queue differ from the *CodedCharSetId* or *Encoding* fields set in the MQMD structure specified on the MQGET call.
- The *Format* field in the MQMD structure associated with the message is not MQFMT_NONE.
- The *BufferLength* specified on the MQGET call is not zero.
- The message data length is not zero.
- The queue manager supports conversion between the *CodedCharSetId* and *Encoding* fields specified in the MQMD structures associated with the message and the MQGET call. See *CodedCharSetId* and *Encoding* for details of the coded character set identifiers and machine encodings supported.

- The queue manager supports conversion of the message format. If the *Format* field of the MQMD structure associated with the message is one of the built-in formats, the queue manager can convert the message. If the *Format* is not one of the built-in formats, you need to write a data-conversion exit to convert the message.

If the sending MCA is to convert the data, specify the CONVERT(YES) keyword on the definition of each sender or server channel for which conversion is required. If the data conversion fails, the message is sent to the DLQ at the sending queue manager and the *Feedback* field of the MQDLH structure indicates the reason. If the message cannot be put on the DLQ, the channel closes and the unconverted message remains on the transmission queue. Data conversion within applications rather than at sending MCAs avoids this situation.

As a rule, data in the message that is described as *character* data by the built-in format or data-conversion exit is converted from the coded character set used by the message to that requested, and *numeric* fields are converted to the encoding requested.

For further details of the conversion processing conventions used when converting the built-in formats, and for information about writing your own data-conversion exits, see “Writing data-conversion exits” on page 364. See also National languages and Machine encodings for information about the language support tables and about the supported machine encodings.

Conversion of EBCDIC newline characters

If you need to ensure that the data that you send from an EBCDIC platform to an ASCII one is identical to the data that you receive back again, you must control the conversion of EBCDIC newline characters.

You can do this using a platform-dependent switch that forces IBM MQ to use the unmodified conversion tables, but you must be aware of the inconsistent behavior that might result.

The problem arises because the EBCDIC newline character is not converted consistently across platforms or conversion tables. As a result, if the data is displayed on an ASCII platform, the formatting might be incorrect. This would make it difficult, for example, to administer an IBM i system remotely from an ASCII platform using RUNMQSC.

See Data conversion for further information about converting EBCDIC-format data to ASCII format.

Browsing messages on a queue:

Use this information to find out about browsing messages on a queue using the MQGET call.

To use the MQGET call to browse the messages on a queue:

1. Call MQOPEN to open the queue for browsing, specifying the MQOO_BROWSE option.
2. To browse the first message on the queue, call MQGET with the MQGMO_BROWSE_FIRST option. To find the message that you want, call MQGET repeatedly with the MQGMO_BROWSE_NEXT option to step through many messages.

You must set the *MsgId* and *CorrelId* fields of the MQMD structure to null after each MQGET call in order to see all messages.

3. Call MQCLOSE to close the queue.

The browse cursor:

When you open (MQOPEN) a queue for browsing, the call establishes a browse cursor for use with MQGET calls that use one of the browse options. You can think of the browse cursor as a logical pointer that is positioned before the first message on the queue.

You can have more than one browse cursor active (from a single program) by issuing several MQOPEN requests for the same queue.

When you call MQGET for browsing, use one of the following options in your MQGMO structure:

MQGMO_BROWSE_FIRST

Gets a copy of the first message that satisfies the conditions specified in your MQMD structure.

MQGMO_BROWSE_NEXT

Gets a copy of the next message that satisfies the conditions specified in your MQMD structure.

MQGMO_BROWSE_MSG_UNDER_CURSOR

Gets a copy of the message currently pointed to by the cursor, that is, the one that was last retrieved using either the MQGMO_BROWSE_FIRST or the MQGMO_BROWSE_NEXT option.

In all cases, the message remains on the queue.

When you open a queue, the browse cursor is positioned logically just before the first message on the queue. This means that if you make your MQGET call immediately after your MQOPEN call, you can use the MQGMO_BROWSE_NEXT option to browse the first message; you do not have to use the MQGMO_BROWSE_FIRST option.

The order in which messages are copied from the queue is determined by the **MsgDeliverySequence** attribute of the queue. (For more information, see “The order in which messages are retrieved from a queue” on page 133.)

- “Queues in FIFO (first in, first out) sequence”
- “Queues in priority sequence”
- “Uncommitted messages” on page 167
- “Change to queue sequence” on page 167
- “Using the queue’s index” on page 167

Queues in FIFO (first in, first out) sequence

The first message in a queue in this sequence is the message that has been on the queue the longest.

Use MQGMO_BROWSE_NEXT to read the messages sequentially in the queue. You will see any messages put to the queue while you are browsing, as a queue in this sequence has messages placed at the end. When the cursor recognizes that it has reached the end of the queue, the browse cursor stays where it is and returns with MQRC_NO_MSG_AVAILABLE. You can then either leave it there waiting for further messages or reset it to the beginning of the queue with a MQGMO_BROWSE_FIRST call.

Queues in priority sequence

The first message in a queue in this sequence is the message that has been on the queue the longest and that has the highest priority at the time that the MQOPEN call is issued.

Use MQGMO_BROWSE_NEXT to read the messages in the queue.

The browse cursor points to the next message, working from the priority of the first message to finish with the message at the lowest priority. It browses any messages put to the queue during this time as long as they are of priority equal to, or lower than, the message identified by the current browse cursor.

Any messages put to the queue of higher priority can be browsed only by:

- Opening the queue for browse again, at which point a new browse cursor is established
- Using the MQGMO_BROWSE_FIRST option

Uncommitted messages

An uncommitted message is never visible to a browse; the browse cursor skips past it.

Messages within a unit-of-work cannot be browsed until the unit-of-work is committed. Messages do not change their position on the queue when committed, so skipped, uncommitted messages will not be seen, even when they *are* committed, unless you use the MQGMO_BROWSE_FIRST option and work through the queue again.

Change to queue sequence

If the message delivery sequence is changed from priority to FIFO while there are messages on the queue, the order of the messages that are already queued is not changed. Messages added to the queue later, take the default priority of the queue.

Using the queue's index

When you browse an indexed queue that contains only messages of a single priority (either persistent or nonpersistent or both), the queue manager uses the index to browse when certain forms of browse are used.

Note: Supported only on IBM MQ for z/OS.

Any of the following forms of browse are used when an indexed queue contains only messages of single priority:

1. If the queue is indexed by MSGID, browse requests that pass a MSGID in the MQMD structure are processed using the index to find the target message.
2. If the queue is indexed by CORRELID, browse requests that pass a CORRELID in the MQMD structure are processed using the index to find the target message.
3. If the queue is indexed by GROUPLID, browse requests that pass a GROUPLID in the MQMD structure are processed using the index to find the target message.

If the browse request does not pass a MSGID, CORRELID, or GROUPLID in the MQMD structure, the queue is indexed, and a message is returned, the index entry for the message must be found, and information within it used to update the browse cursor. If you use a wide selection of index values, this does not add significant extra processing to the browse request.

Browsing messages when the message length is unknown:

To browse a message when you do not know the size of the message, and you do not want to use the *MsgId*, *CorrelId*, or *GroupId* fields to locate the message, you can use the `MQGMO_BROWSE_MSG_UNDER_CURSOR` option:

1. Issue an `MQGET` with:
 - Either the `MQGMO_BROWSE_FIRST` or `MQGMO_BROWSE_NEXT` option
 - The `MQGMO_ACCEPT_TRUNCATED_MSG` option
 - Buffer length zero

Note: If another program is likely to get the same message, consider using the `MQGMO_LOCK` option as well. `MQRCTRUNCATEDMSGACCEPTED` should be returned.

2. Use the returned *DataLength* to allocate the storage needed.
3. Issue an `MQGET` with the `MQGMO_BROWSE_MSG_UNDER_CURSOR`.

The message pointed to is the last one that was retrieved; the browse cursor will not have moved. You can choose either to lock the message using the `MQGMO_LOCK` option, or to unlock a locked message using `MQGMO_UNLOCK` option.

The call fails if no `MQGET` with either the `MQGMO_BROWSE_FIRST` or `MQGMO_BROWSE_NEXT` options has been issued successfully since the queue was opened.

Removing a message that you have browsed:

You can remove from the queue a message that you have already browsed provided that you have opened the queue for removing messages as well as for browsing. (You must specify one of the `MQOO_INPUT_*` options, as well as the `MQOO_BROWSE` option, on your `MQOPEN` call.)

To remove the message, call `MQGET` again, but in the *Options* field of the `MQGMO` structure, specify `MQGMO_MSG_UNDER_CURSOR`. In this case, the `MQGET` call ignores the *MsgId*, *CorrelId*, and *GroupId* fields of the `MQMD` structure.

In the time between your browsing and removal steps, another program might have removed messages from the queue, including the message under your browse cursor. In this case, your `MQGET` call returns a reason code to say that the message is not available.

Browsing messages in logical order:

“Logical and physical ordering” on page 134 explains the difference between the logical and physical order of messages on a queue. This distinction is particularly important when browsing a queue, because, in general, messages are not being deleted and browse operations do not necessarily start at the beginning of the queue.

If an application browses through the various messages of one group (using logical order), it is important that logical order should be followed to reach the start of the next group, because the last message of one group might occur physically *after* the first message of the next group. The `MQGMO_LOGICAL_ORDER` option ensures that logical order is followed when scanning a queue.

Use `MQGMO_ALL_MSGS_AVAILABLE` (or `MQGMO_ALL_SEGMENTS_AVAILABLE`) with care for browse operations. Consider the case of logical messages with `MQGMO_ALL_MSGS_AVAILABLE`. The effect of this is that a logical message is available only if all the remaining messages in the group are also present. If they are not, the message is passed over. This can mean that when the missing messages arrive subsequently, they are not noticed by a browse-next operation.

For example, if the following logical messages are present,


```
Logical message 1 (not last) of group 123
Logical message 1 (not last) of group 456
Logical message 2 (last) of group 456
```

and a browse function is issued with `MQGMO_ALL_MSGS_AVAILABLE`, the first logical message of group 456 is returned, leaving the browse cursor on this logical message. If the second (last) message of group 123 now arrives:

```
Logical message 1 (not last) of group 123
Logical message 2 (last) of group 123
Logical message 1 (not last) of group 456 <=== browse cursor
Logical message 2 (last) of group 456
```

and the same browse-next function is issued, it is not noticed that group 123 is now complete, because the first message of this group is *before* the browse cursor.

In some cases (for example, if messages are retrieved destructively when the group is present in its entirety), you can use `MQGMO_ALL_MSGS_AVAILABLE` together with `MQGMO_BROWSE_FIRST`. Otherwise, you must repeat the browse scan to take note of newly-arrived messages that have been missed; just issuing `MQGMO_WAIT` together with `MQGMO_BROWSE_NEXT` and `MQGMO_ALL_MSGS_AVAILABLE` does not take account of them. (This also happens to higher-priority messages that might arrive after scanning the messages is complete.)

The next sections look at browsing examples that deal with unsegmented messages; segmented messages follow similar principles.

Browsing messages in groups:

In this example, the application browses through each message on the queue, in logical order.

Messages on the queue might be grouped. For grouped messages, the application does not want to start processing any group until all the messages within it have arrived. `MQGMO_ALL_MSGS_AVAILABLE` is therefore specified for the first message in the group; for subsequent messages in the group, this option is unnecessary.

`MQGMO_WAIT` is used in this example. However, although the wait can be satisfied if a new group arrives, for the reasons in “Browsing messages in logical order” on page 168, it is not satisfied if the browse cursor has already passed the first logical message in a group, and the remaining messages now arrive. Nevertheless, waiting for a suitable interval ensures that the application does not constantly loop while waiting for new messages or segments.

`MQGMO_LOGICAL_ORDER` is used throughout, to ensure that the scan is in logical order. This contrasts with the destructive `MQGET` example, where because each group is being removed, `MQGMO_LOGICAL_ORDER` is not used when looking for the first (or only) message in a group.

It is assumed that the application's buffer is always large enough to hold the entire message, whether or not the message has been segmented. `MQGMO_COMPLETE_MSG` is therefore specified on each `MQGET`.

The following gives an example of browsing logical messages in a group:

```
/* Browse the first message in a group, or a message not in a group */
GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
| MQGMO_ALL_MSGS_AVAILABLE | MQGMO_WAIT
MQGET GMO.MatchOptions = MQMO_MATCH_MSG_SEQ_NUMBER, MD.MsgSeqNumber = 1
/* Examine first or only message */
...
GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
```

```

do while ( GroupStatus == MQGS_MSG_IN_GROUP )
  MQGET
  /* Examine each remaining message in the group */
  ...

```

The group is repeated until MQRC_NO_MSG_AVAILABLE is returned.

Browsing and retrieving destructively:

In this example, the application browses each of the logical messages within a group, before deciding whether to retrieve that group destructively.

The first part of this example is similar to the previous one. However, in this case, having browsed an entire group, we decide to go back and retrieve it destructively.

As each group is removed in this example, MQGMO_LOGICAL_ORDER is not used when looking for the first or only message in a group.

The following gives an example of browsing and then retrieving destructively:

```

GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
              | MQGMO_ALL_MESSAGES_AVAILABLE | MQGMO_WAIT
do while ( GroupStatus == MQGS_MSG_IN_GROUP )
  MQGET
  /* Examine each remaining message in the group (or as many as
     necessary to decide whether to get it destructively) */
  ...

if ( we want to retrieve the group destructively )

  if ( GroupStatus == ' ' )
    /* We retrieved an ungrouped message */
    GMO.Options = MQGMO_MSG_UNDER_CURSOR | MQGMO_SYNCPOINT
    MQGET GMO.MatchOptions = 0
    /* Process the message */
    ...

  else
    /* We retrieved one or more messages in a group. The browse cursor */
    /* will not normally be still on the first in the group, so we have */
    /* to match on the GroupId and MsgSeqNumber = 1. */
    /* Another way, which works for both grouped and ungrouped messages, */
    /* would be to remember the MsgId of the first message when it was */
    /* browsed, and match on that. */
    GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT
    MQGET GMO.MatchOptions = MQMO_MATCH_GROUP_ID
                      | MQMO_MATCH_MSG_SEQ_NUMBER,
          (MQMD.GroupId      = value already in the MD)
          MQMD.MsgSeqNumber = 1
    /* Process first or only message */
    ...

  GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT
                | MQGMO_LOGICAL_ORDER
do while ( GroupStatus == MQGS_MSG_IN_GROUP )
  MQGET
  /* Process each remaining message in the group */
  ...

```

Avoiding repeated delivery of browsed messages:

By using certain open options and get-message options, you can mark messages as having been browsed so that they are not retrieved again by the current or other cooperating applications. Messages can be unmarked explicitly or automatically to make them available again for browsing.

If you browse messages on a queue, you might retrieve them in a different order to the order in which you would retrieve them if you got them destructively. In particular, you can browse the same message multiple times, which is not possible if it is removed from the queue. To avoid this you can *mark* messages as they are browsed, and avoid retrieving marked messages. This is sometimes referred to as *browse with mark*. To mark browsed messages, use the get message option `MQGMO_MARK_BROWSE_HANDLE`, and to retrieve only messages that are not marked, use `MQGMO_UNMARKED_BROWSE_MSG`. If you use the combination of options `MQGMO_BROWSE_FIRST`, `MQGMO_UNMARKED_BROWSE_MSG`, and `MQGMO_MARK_BROWSE_HANDLE`, and issue repeated `MQGET`s, you will retrieve each message on the queue in turn. This prevents repeated delivery of messages even though `MQGMO_BROWSE_FIRST` is used to ensure that messages are not skipped. This combination of options can be represented by the single constant `MQGMO_BROWSE_HANDLE`. When there are no messages on the queue that have not been browsed, `MQRC_NO_MSG_AVAILABLE` is returned.

If multiple applications are browsing the same queue, they can open the queue with the options `MQOO_CO_OP` and `MQOO_BROWSE`. The object handle returned by each `MQOPEN` is considered to be part of a cooperating group. Any message returned by an `MQGET` call specifying the option `MQGMO_MARK_BROWSE_CO_OP` is considered to be marked for this cooperating set of handles.

If a message has been marked for some time, it can be automatically unmarked by the queue manager and made available for browsing again. The queue manager attribute `MsgMarkBrowseInterval` gives the time in milliseconds for which a message is to remain marked for the cooperating set of handles. A `MsgMarkBrowseInterval` of -1 means that messages are never automatically unmarked.

When the single process or set of cooperative processes marking messages stop, any marked messages become unmarked.

Examples of cooperative browsing

You might run multiple copies of a dispatcher application to browse messages on a queue and initiate a consumer based on the content of each message. In each dispatcher, open the queue with `MQOO_CO_OP`. This indicates that the dispatchers are cooperating and will be aware of each other's marked messages. Each dispatcher then makes repeated `MQGET` calls, specifying the options `MQGMO_BROWSE_FIRST`, `MQGMO_UNMARKED_BROWSE_MSG`, and `MQGMO_MARK_BROWSE_CO_OP` (you can use the single constant `MQGMO_BROWSE_CO_OP` to represent this combination of options). Each dispatcher application then retrieves only those messages that have not already been marked by other cooperating dispatchers. The dispatcher initializes a consumer and passes the `MsgToken` returned by the `MQGET` to the consumer, which destructively gets the message from the queue. If the consumer backs out the `MQGET` of the message, then the message is available for one of the browsers to re-dispatch, because it is no longer marked. If the consumer does not do an `MQGET` on the message, then after the `MsgMarkBrowseInterval` has passed, the queue manager unmarks the message for the cooperating set of handles, and it can be re-dispatched.

Rather than multiple copies of the same dispatcher application, you might have a number of different dispatcher applications browsing the queue, each suitable for processing a subset of the messages on the queue. In each dispatcher, open the queue with `MQOO_CO_OP`. This indicates that the dispatchers are cooperating and will be aware of each other's marked messages.

- If the order of message processing for a single dispatcher is important, each dispatcher makes repeated `MQGET` calls, specifying the options `MQGMO_BROWSE_FIRST`, `MQGMO_UNMARKED_BROWSE_MSG`, and `MQGMO_MARK_BROWSE_HANDLE` (or

MQGMO_BROWSE_HANDLE). If the browsed message is suitable for this dispatcher to process, it then makes an MQGET call specifying MQMO_MATCH_MSG_TOKEN, MQGMO_MARK_BROWSE_CO_OP, and the MsgToken returned by the previous MQGET call. If the call succeeds, the dispatcher initializes the consumer, passing the MsgToken to it.

- If the order of message processing is not important and the dispatcher is expected to process most of the messages it encounters, use the options MQGMO_BROWSE_FIRST, MQGMO_UNMARKED_BROWSE_MSG, and MQGMO_MARK_BROWSE_CO_OP (or MQGMO_BROWSE_CO_OP). If the dispatcher browses a message it cannot process, it unmarks the message by calling MQGET with the option MQMO_MATCH_MSG_TOKEN, MQGMO_UNMARK_BROWSE_CO_OP, and the MsgToken returned previously.

Some cases where the MQGET call fails:

If certain attributes of a queue are changed using the FORCE option on a command between issuing an MQOPEN and an MQGET call, the MQGET call fails and returns the MQRC_OBJECT_CHANGED reason code.

The queue manager marks the object handle as being no longer valid. This also happens if the changes apply to any queue to which the queue name resolves. The attributes that affect the handle in this way are listed in the description of the MQOPEN call in MQOPEN. If your call returns the MQRC_OBJECT_CHANGED reason code, close the queue, reopen it, then try to get a message again.

If get operations are inhibited for a queue from which you are attempting to get messages (or any queue to which the queue name resolves), the MQGET call fails and returns the MQRC_GET_INHIBITED reason code. This happens even if you are using the MQGET call for browsing. You might be able to get a message successfully if you attempt the MQGET call at a later time, if the design of the application is such that other programs change the attributes of queues regularly.

If a dynamic queue (either temporary or permanent) has been deleted, MQGET calls using a previously-acquired object handle fail and return the MQRC_Q_DELETED reason code.

Writing publish/subscribe applications

Start writing publish/subscribe IBM MQ applications.

For an overview of publish/subscribe concepts, see Publish/subscribe messaging.

See the following topics for information on writing different types of publish/subscribe applications:

- “Writing publisher applications” on page 173
- “Writing subscriber applications” on page 181
- “Publish/subscribe lifecycles” on page 201
- “Publish/subscribe message properties” on page 206
- “Message ordering” on page 208
- “Intercepting publications” on page 208
- “Publishing options” on page 216
- “Subscription options” on page 217

Related concepts:

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

“Developing applications for IBM MQ” on page 1

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM MQ support applications written in procedural languages, and object oriented languages and frameworks.

“Design considerations for IBM MQ applications” on page 43

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

“Writing a procedural application for queuing” on page 74

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

“Writing client procedural applications” on page 285

What you need to know to write client applications on IBM MQ using a procedural language.

“Building a procedural application” on page 382

You can write an IBM MQ application in one of several procedural languages, and run the application on several different platforms.

“Handling procedural program errors” on page 436

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Related tasks:

“Using the IBM MQ sample procedural programs” on page 457

These sample programs are written in procedural languages, and demonstrate typical uses of the Message Queue Interface (MQI). IBM MQ programs on different platforms.

“Developing web services with IBM MQ” on page 1217

You can develop IBM MQ applications for web services using the IBM MQ transport for SOAP.

Writing publisher applications:

Get started with writing publisher applications by studying two examples. The first is modeled as closely as possible on a point to point application putting messages on a queue, and the second demonstrates creating topics dynamically - a more common pattern for publisher applications.

Writing a simple IBM MQ publisher application is just like writing an IBM MQ point to point application that puts messages to a queue (Table 23). The difference is you MQPUT messages to a topic, not to a queue.

Table 23. Point to point versus publish/subscribe IBM MQ program pattern.

Step	Point to point MQ Call	Publish MQ Call
Connect to a queue manager	MQCONN	MQCONN
Open queue	MQOPEN	
Open topic		MQOPEN
Put message(s)	MQPUT	MQPUT
Close topic		MQCLOSE
Close queue	MQCLOSE	
Disconnect from queue manager	MQDISC	MQDISC

To make that concrete, there are two examples of applications to publish stock prices. In the first example (“Example 1: Publisher to a fixed topic” on page 174), that is modeled closely on putting messages to a queue, the administrator creates a topic definition in a similar way to creating a queue. The programmer

codes MQPUT to write messages to the topic instead of writing them to a queue. In the second example (“Example 2: Publisher to a variable topic” on page 178), the pattern of interaction of the program with IBM MQ is similar. The difference is the programmer provides the topic to which the message is written, rather than the administrator. In practice this typically means that the topic string is content defined, or provided by another source, such as human input through a browser.

Related concepts:

“Writing subscriber applications” on page 181

Get started with writing subscriber applications by studying three examples: an IBM MQ application consuming messages from a queue, an application that creates a subscription and requires no knowledge of queuing, and finally an example that uses both queuing and subscriptions.

Related information:

DEFINE TOPIC

DISPLAY TOPIC

DISPLAY TPSTATUS

Example 1: Publisher to a fixed topic:

An IBM MQ program to illustrate publishing to an administratively defined topic.

Note: The compact coding style is intended for readability not production use.

See the output in Figure 18 on page 176

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>
int main(int argc, char **argv)
{
    char    topicNameDefault[] = "IBMSTOCKPRICE";
    char    publicationDefault[] = "129";
    MQCHAR48 qmName = "";

    MQHCONN Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ  Hobj = MQHO_NONE;           /* object handle sub queue */
    MQLONG  CompCode = MQCC_OK;         /* completion code */
    MQLONG  Reason = MQRC_NONE;        /* reason code */
    MQOD    td = {MQOD_DEFAULT};       /* Object descriptor */
    MQMD    md = {MQMD_DEFAULT};       /* Message Descriptor */
    MQPMO   pmo = {MQPMO_DEFAULT};     /* put message options */
    MQCHAR  resTopicStr[151];         /* Returned vale of topic string */
    char *  topicName = topicNameDefault;
    char *  publication = publicationDefault;
    memset (resTopicStr, 0 , sizeof(resTopicStr));

    switch(argc){                      /* replace defaults with args if provided */
        default:
            publication = argv[2];
        case(2):
            topicName = argv[1];
        case(1):
            printf("Optional parameters: TopicObject Publication\n");
    }
    do {
        MQCONN(qmName, &Hconn, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        td.ObjectType = MQOT_TOPIC;     /* Object is a topic */
        td.Version = MQOD_VERSION_4;    /* Descriptor needs to be V4 */
        strncpy(td.ObjectName, topicName, MQ_TOPIC_NAME_LENGTH);
        td.ResObjectString.VSPtr = resTopicStr;
        td.ResObjectString.VSBufSize = sizeof(resTopicStr)-1;
        MQOPEN(Hconn, &td, MQOO_OUTPUT | MQOO_FAIL_IF QUIESCING, &Hobj, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        pmo.Options = MQPMO_FAIL_IF QUIESCING | MQPMO_RETAIN;
        MQPUT(Hconn, Hobj, &md, &pmo, (MQLONG)strlen(publication)+1, publication, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        MQDISC(&Hconn, &CompCode, &Reason);
    } while (0);
    if (CompCode == MQCC_OK)
        printf("Published \"%s\" using topic \"%s\" to topic string \"%s\"\n",
            publication, td.ObjectName, resTopicStr);
    printf("Completion code %d and Return code %d\n", CompCode, Reason);
}
```

Figure 17. Simple IBM MQ publisher to a fixed topic.

```

X:\Publish1\Debug>PublishStock
Optional parameters: TopicObject Publication
Published "129" using topic "IBMSTOCKPRICE" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

X:\Publish1\Debug>PublishStock IBMSTOCKPRICE 155
Optional parameters: TopicObject Publication
Published "155" using topic "IBMSTOCKPRICE" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

```

Figure 18. Sample output from first publisher example

The following selected lines of code illustrate aspects of writing a publisher application for IBM MQ.

char topicNameDefault[] = "IBMSTOCKPRICE";

A default topic name is defined in the program. You can override it by providing the name of a different topic object as the first argument to the program.

MQCHAR resTopicStr[151];

resTopicStr is pointed at by td.ResObjectString.VSPtr and is used by MQOPEN to return the resolved topic string. Make the length of resTopicStr one larger than the length passed in td.ResObjectString.VSBufSize to give space for null termination.

memset (resTopicStr, 0, sizeof(resTopicStr));

Initialize resTopicStr to nulls to ensure the resolved topic string returned in an MQCHARV is null terminated.

td.ObjectType = MQOT_TOPIC

There is a new type of object for publish/subscribe: the *topic object*.

td.Version = MQOD_VERSION_4;

To use the new type of object, you must use at least *version 4* of the object descriptor.

strncpy(td.ObjectName, topicName, MQ_OBJECT_NAME_LENGTH);

The topicName is the name of a topic object, sometimes called an administrative topic object. In the example the topic object needs to be created beforehand, using IBM MQ Explorer or this MQSC command,

```
DEFINE TOPIC(IBMSTOCKPRICE) TOPICSTR(NYSE/IBM/PRICE) REPLACE;
```

td.ResObjectString.VSPtr = resTopicStr;

The resolved topic string is echoed in the final printf in the program. Set up the MQCHARV ResObjectString structure for IBM MQ to return the resolved string back to the program.

MQOPEN(Hconn, &td, MQOO_OUTPUT | MQOO_FAIL_IF QUIESCING, &Hobj, &CompCode, &Reason);

Open the topic for output; just like opening a queue for output.

pmo.Options = MQPMO_FAIL_IF QUIESCING | MQPMO_RETAIN;

You want new subscribers to be able receive the publication, and by specifying MQPMO_RETAIN in the publisher, when you start a subscriber it receives the latest publication, published before the subscriber started, as its first matching publication. The alternative is to provide subscribers with publications published only after the subscriber started. Additionally a subscriber has the option to decline to receive a retained publication by specifying MQSO_NEW_PUBLICATIONS_ONLY in its subscription.

MQPUT(Hconn, Hobj, &md, &pmo, (MQLONG)strlen(publication)+1, publication, &CompCode, &Reason);

Add 1 to the length of the string passed to MQPUT to pass the null termination character to IBM MQ as part of the message buffer.

What does the first example demonstrate? The example imitates as closely as possible the tried and tested traditional pattern for writing point to point IBM MQ programs. An important feature of the IBM MQ programming pattern is that the programmer is not concerned where messages are sent. The task of the programmer is to connect to a queue manager, and pass it the messages that are to be distributed to recipients. In the point-to-point paradigm, the programmer opens a queue (probably an alias queue) that

the administrator has configured. The alias queue routes messages to a target queue, either on the local queue manager, or to a remote queue manager. While the messages are waiting to be delivered, they are stored on queues somewhere between the source and the destination.

In the publish/subscribe pattern, instead of opening a queue, the programmer opens a topic. In our example, the topic is associated with a topic string by an administrator. The queue manager forwards the publication, using queues, to local or remote subscribers that have subscriptions that match the topic string of the publication. If publications are retained the queue manager keeps the latest copy of the publication, even if it has no subscribers now. The retained publication is available to forward to future subscribers. The publisher application plays no part in selecting or routing the publication to a destination; its task is to create and put publications to the topics defined by the administrator.

This fixed topic example is atypical of many publish/subscribe applications: it is static. It requires an administrator to define the topic strings and change the topics that are published on. Commonly publish/subscribe applications need to know some or all the topic tree. Perhaps topics change frequently, or perhaps although the topics do not change much, the number of topic combinations is large and it is too onerous for an administrator to define a topic node for every topic string that might need to be published on. Perhaps topic strings are not known in advance of publication; a publisher application might use information from the publication content to specify a topic string, or it might have information about topic strings to publish on from another source, such as human input from a browser. To cater for more dynamic styles of publishing, the next example shows how to create topics dynamically, as part of the publisher application.

Topics couple publishers and subscribers together. Designing the rules, or architecture, for naming topics, and organizing them in topic trees is an important step in developing a publish/subscribe solution. Look carefully at the extent to which organization of the topic tree binds of publisher and subscriber programs together, and binds them to the content of the topic tree. Ask yourself the question whether changes in the topic tree affect publisher and subscriber applications, and how you can minimize the effect. Built into the architecture of the IBM MQ publish/subscribe model is the notion of an administrative topic object that provides the root part, or root subtree, of a topic. The topic object gives you the option of defining the root part of the topic tree administratively that simplifies application programming and operations, and consequently improves maintainability. For example, if you are deploying multiple publish/subscribe applications that have isolated topic trees, then by administratively defining the root part of the topic tree, you can guarantee the isolation of topic trees, even if there is no consistency in the topic naming conventions adopted by the different applications.

In practice, publisher applications cover a spectrum from solely using fixed topics, as in this example, and variable topics, as in the next. “Example 2: Publisher to a variable topic” on page 178 also demonstrates combining the use of topics and topic strings.

Related concepts:

“Example 2: Publisher to a variable topic”

A WebSphere MQ program to illustrate publishing to a programmatically defined topic.

“Writing subscriber applications” on page 181

Get started with writing subscriber applications by studying three examples: an IBM MQ application consuming messages from a queue, an application that creates a subscription and requires no knowledge of queuing, and finally an example that uses both queuing and subscriptions.

Example 2: Publisher to a variable topic:

A WebSphere MQ program to illustrate publishing to a programmatically defined topic.

Note: The compact coding style is intended for readability not production use.

See the output in Figure 20 on page 180.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>
int main(int argc, char **argv)
{
    char    topicNameDefault[] = "STOCKS";
    char    topicStringDefault[] = "IBM/PRICE";
    char    publicationDefault[] = "130";
    MQCHAR48 qmName = "";

    MQHCONN Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ  Hobj  = MQHO_NONE;           /* object handle sub queue */
    MQLONG  CompCode = MQCC_OK;         /* completion code */
    MQLONG  Reason  = MQRC_NONE;        /* reason code */
    MQOD    td = {MQOD_DEFAULT};       /* Object descriptor */
    MQMD    md = {MQMD_DEFAULT};       /* Message Descriptor */
    MQPMO   pmo = {MQPMO_DEFAULT};     /* put message options */
    MQCHAR  resTopicStr[151];          /* Returned value of topic string */
    char *  topicName = topicNameDefault;
    char *  topicString = topicStringDefault;
    char *  publication = publicationDefault;
    memset (resTopicStr, 0 , sizeof(resTopicStr));

    switch(argc){ /* Replace defaults with args if provided */
    default:
        publication = argv[3];
    case(3):
        topicString = argv[2];
    case(2):
        if (strcmp(argv[1],"/")) /* "/" invalid = No topic object */
            topicName = argv[1];
        else
            *topicName = '\0';
    case(1):
        printf("Provide parameters: TopicObject TopicString Publication\n");
    }

    printf("Publish \"%s\" to topic \"%-.48s\" and topic string \"%s\"\n", publication, topicName, topicString);
    do {
        MQCONN(qmName, &Hconn, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        td.ObjectType = MQOT_TOPIC; /* Object is a topic */
        td.Version = MQOD_VERSION_4; /* Descriptor needs to be V4 */
        strncpy(td.ObjectName, topicName, MQ_TOPIC_NAME_LENGTH);
        td.ObjectString.VSPtr = topicString;
        td.ObjectString.VSLength = (MQLONG)strlen(topicString);
        td.ResObjectString.VSPtr = resTopicStr;
        td.ResObjectString.VSBufSize = sizeof(resTopicStr)-1;
        MQOPEN(Hconn, &td, MQOO_OUTPUT | MQOO_FAIL_IF QUIESCING, &Hobj, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        pmo.Options = MQPMO_FAIL_IF QUIESCING | MQPMO_RETAIN;
        MQPUT(Hconn, Hobj, &md, &pmo, (MQLONG)strlen(publication)+1, publication, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        MQDISC(&Hconn, &CompCode, &Reason);
    } while (0);
    if (CompCode == MQCC_OK)
        printf("Published \"%s\" to topic string \"%s\"\n", publication, resTopicStr);
    printf("Completion code %d and Return code %d\n", CompCode, Reason);
}
```

Figure 19. Simple IBM MQ publisher to a variable topic.

```

X:\Publish2\Debug>PublishStock
Provide parameters: TopicObject TopicString Publication
Publish "130" to topic "STOCKS" and topic string "IBM/PRICE"
Published "130" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

X:\Publish2\Debug>PublishStock / NYSE/IBM/PRICE 131
Provide parameters: TopicObject TopicString Publication
Publish "131" to topic "" and topic string "NYSE/IBM/PRICE"
Published "131" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

```

Figure 20. Sample output from second publisher example

There are a few points to note about this example.

```
char topicNameDefault[] = "STOCKS";
```

The default topic name STOCKS defines part of the topic string. You can override this topic name by providing it as the first argument to the program, or eliminate the use of the topic name by supplying / as the first parameter.

```
char topicString[101] = "IBM/PRICE";
```

IBM/PRICE is the default topic string. You can override this topic string by providing it as the second argument to the program.

The queue manager combines the topic string provided by the STOCKS topic object, "NYSE", with the topic string provided by the program "IBM/PRICE" and inserts a "/" between the two topic strings. The result is the resolved topic string "NYSE/IBM/PRICE". The resulting topic string is the same as the one defined in the IBMSTOCKPRICE topic object, and has precisely the same effect.

The administrative topic object associated with the resolved topic string is not necessarily the same topic object as passed to MQOPEN by the publisher. IBM MQ uses the tree implicit in the resolved topic string to work out which administrative topic object defines the attributes associated with the publication.

Suppose there are two topic objects A and B, and A defines topic "a", and B defines topic "a/b" (Figure 21 on page 181). If the publisher program refers to topic object A and provides topic string "b", resolving the topic to the topic string "a/b", then the publication inherits its properties from topic object B because the topic matches the topic string "a/b" defined for B.

```
if (strcmp(argv[1],"/"))
```

argv[1] is the optionally provided topicName. "/" is invalid as a topic name; here it signifies that there is no topic name, and the topic string is provided entirely by the program. The output in Figure 20 shows the whole topic string being supplied dynamically by the program.

```
strncpy(td.ObjectName, topicName, MQ_OBJECT_NAME_LENGTH);
```

For the default case, the optional topicName needs to be created beforehand, using IBM MQ Explorer or this MQSC command:

```
DEFINE TOPIC(STOCKS) TOPICSTR(NYSE) REPLACE;
```

```
td.ObjectString.VSPtr = topicString;
```

The topic string is a MQCHARV field in the topic descriptor



Figure 21. Topic object associations

What does the second example demonstrate? Although the code is very similar to the first example - effectively there are only two lines difference - the result is a significantly different program to the first. The programmer controls the destinations to which publications are sent. In conjunction with minimal administrator input used to design subscriber applications, no topics or queues need to be predefined to route publications from publishers to subscribers.

In the point-to-point messaging paradigm, queues have to be defined before messages are able to flow. For publish/subscribe, they do not, although IBM MQ implements publish/subscribe using its underlying queuing system; the benefits of guaranteed delivery, transactionality and loose coupling associated with messaging and queuing are inherited by publish/subscribe applications.

A designer has to decide whether publisher, and subscriber, programs are to be aware of the underlying topic tree or not, and also whether subscriber programs are aware of queuing or not. Study the subscriber example applications next. They are designed to be used with the publisher examples, typically publishing and subscribing to NYSE/IBM/PRICE.

Related concepts:

“Example 1: Publisher to a fixed topic” on page 174

An IBM MQ program to illustrate publishing to an administratively defined topic.

“Writing subscriber applications”

Get started with writing subscriber applications by studying three examples: an IBM MQ application consuming messages from a queue, an application that creates a subscription and requires no knowledge of queuing, and finally an example that uses both queuing and subscriptions.

Writing subscriber applications:

Get started with writing subscriber applications by studying three examples: an IBM MQ application consuming messages from a queue, an application that creates a subscription and requires no knowledge of queuing, and finally an example that uses both queuing and subscriptions.

In Table 24 on page 182 the three styles of consumer or subscriber are listed, together with the sequences of IBM MQ function calls that characterize them.

1. The first style, MQ Publication Consumer, is identical to a point to point MQ program that only does MQGET. The application has no knowledge that it is consuming publications - it is simply reading messages from a queue. The subscription that causes publications to get routed to the queue is created administratively using IBM MQ Explorer or a command.
2. The second style is the preferred pattern for most subscriber applications. The subscriber application creates the subscription, and then gets publications. The queue management is all performed by the queue manager.

3. In the third style, the subscriber application elects to open and close the underlying queue that is used for publications as well as issue subscriptions to fill the queue with publications.

One way to understand these styles is to study the example C programs listed in Table 24 for each of the styles. The examples are designed to be run in conjunction with the publisher example found in “Writing publisher applications” on page 173.

Table 24. Point to point vs. subscribe IBM MQ program patterns.

Step	MQ message consumer	“Example 1: MQ Publication consumer”	“Example 2: Managed MQ subscriber” on page 185	“Example 3: Unmanaged MQ subscriber” on page 191
Connect to a queue manager	MQCONN	MQCONN	MQCONN	MQCONN
Open queue	MQOPEN	MQOPEN		MQOPEN
Subscribe			MQSUB	MQSUB
Get message(s)	MQGET	MQGET	MQGET	MQGET
Close queue	MQCLOSE	MQCLOSE	(MQCLOSE)	MQCLOSE
Close subscription			MQCLOSE	MQCLOSE
Disconnect from queue manager	MQDISC	MQDISC	MQDISC	MQDISC

Using MQCLOSE is always optional, either to release resources, pass MQCLOSE options, or just for symmetry with MQOPEN. Since you are unlikely to need to specify the MQCLOSE options when the subscription queue is closed in the Managed MQ subscriber case, and the symmetry argument is not relevant, the subscription queue is not explicitly closed in Example 2: Managed MQ subscriber.

Another way to understand publish/subscribe application patterns is to look at the interactions between the different entities involved. Lifeline, or UML sequence diagrams are a good way to study interactions. Three lifeline examples are described in “Publish/subscribe lifecycles” on page 201.

Example 1: MQ Publication consumer:

The MQ Publication consumer is an IBM MQ message consumer that does not subscribe to topics itself.

To create the subscription and publication queue for this example run the following commands, or define the objects using IBM MQ Explorer.

```
DEFINE QLOCAL(STOCKTICKER) REPLACE;
DEFINE SUB(IBMSTOCKPRICESUB) DEST(STOCKTICKER) TOPICOBJ(IBMSTOCKPRICE) REPLACE;
```

The IBMSTOCKPRICESUB subscription references the IBMSTOCK topic object created for the publisher example and the local queue STOCKTICKER. The topic object IBMSTOCK defines the topic string that is used in the subscription, NYSE/IBM/PRICE. Note that the topic object and the queue used to receive publications need to be defined before the subscription is created.

There are a number of valuable facets to the MQ publication consumer pattern:

1. Multiprocessing: sharing out of the work of reading publications. The publications all go onto the single queue associated with the subscription topic. Multiple consumers can open the queue using MQ00_INPUT_SHARED.
2. Centrally managed subscriptions. Applications do not construct their own subscription topics or subscriptions; the administrator is responsible for where publications are sent.
3. Subscription concentration: multiple different subscriptions can be sent to a single queue.
4. Subscription durability: the queue receives all publications whether or not consumers are active.

5. Migration and coexistence: the consumer code works equally well for a point-to-point and a publish/subscribe scenario.

The subscription creates a relationship between the topic string NYSE/IBM/PRICE and the queue STOCKTICKER. Publications, including any currently retained publication, are forwarded to STOCKTICKER from the moment the subscription is created.

An administratively created subscription can be managed or unmanaged. A managed subscription takes effect as soon as it has been created, just like an unmanaged subscription. Not all the pattern facets are available to a managed subscription. See “Example 3: Unmanaged MQ subscriber” on page 191

Note: The compact coding style is intended for readability not production use.

The results are shown in Figure 23 on page 185.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>
int main(int argc, char **argv)
{
MQCHAR publicationBuffer[101];
MQCHAR48 subscriptionQueueDefault = "STOCKTICKER";
MQCHAR48 qmName = "";          /* Use default queue manager */

MQHCONN Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
MQHOBJ Hobj = MQHO_NONE; /* object handle sub queue */
MQLONG CompCode = MQCC_OK; /* completion code */
MQLONG Reason = MQRC_NONE; /* reason code */
MQLONG messlen = 0;
MQOD od = {MQOD_DEFAULT}; /* Unmanaged subscription queue */
MQMD md = {MQMD_DEFAULT}; /* Message Descriptor */
MQGMO gmo = {MQGMO_DEFAULT}; /* Get message options */
char * publication=publicationBuffer;
char * subscriptionQueue = subscriptionQueueDefault;

switch(argc){ /* Replace defaults with args if provided */
default:
subscriptionQueue = argv[1]
case(1):
printf("Optional parameter: subscriptionQueue\n");
}

do {
MQCONN(qmName, &Hconn, &CompCode, &Reason);
if (CompCode != MQCC_OK) break;
strncpy(od.ObjectName, subscriptionQueue, MQ_Q_NAME_LENGTH);
MQOPEN(Hconn, &od, MQOO_INPUT_AS_Q_DEF | MQOO_FAIL_IF QUIESCING , &Hobj, &CompCode, &Reason);
if (CompCode != MQCC_OK) break;
gmo.Options = MQGMO_WAIT | MQGMO_NO_SYNCPOINT | MQGMO_CONVERT;
gmo.WaitInterval = 10000;
printf("Waiting %d seconds for publications from %s\n", gmo.WaitInterval/1000, subscriptionQueue);
do {
memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
md.Encoding = MQENC_NATIVE;
md.CodedCharSetId = MQCCSI_Q_MGR;
memset(publication, 0, sizeof(publicationBuffer));
MQGET(Hconn, Hobj, &md, &gmo, sizeof(publicationBuffer)-1, publication, &messlen,
&CompCode, &Reason);
if (Reason == MQRC_NONE)
printf("Received publication \"%s\"\n", publication);
}
while (CompCode == MQCC_OK);
if (CompCode != MQCC_OK && Reason != MQRC_NO_MSG_AVAILABLE) break;
MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);
if (CompCode != MQCC_OK) break;
MQDISC(&Hconn, &CompCode, &Reason);
} while (0);
printf("Completion code %d and Return code %d\n", CompCode, Reason);
}
```

Figure 22. MQ publication consumer.


```
X:\Subscribe1\Debug>Subscribe1
Optional parameter: subscriptionQueue
Waiting 10 seconds for publications from STOCKTICKER
Received publication "129"
Completion code 0 and Return code 0
```

Figure 23. Output from MQ publication consumer

There are a couple of standard IBM MQ C language programming tips to be aware of:

memset(publication, 0, sizeof(publicationBuffer));

Ensure the message has a trailing null for easy formatting using printf. The publisher example includes the trailing null in the message buffer passed to MQPUT by adding 1 to strlen(publication). Setting MQCHAR buffers to null is good programming style for IBM MQ C programs that use the buffers to store strings, ensuring a null follows an array of characters that does not completely fill the buffer.

MQGET(Hconn, Hobj, &md, &gmo, sizeof(publicationBuffer)-1, publication, &messlen, &CompCode, &Reason);

Reserve one null at the end of the message buffer to ensure the returned message has trailing null in case if (messlen == strlen(publication)); is true. This tip complements the preceding one, and ensures that there is at least one null in publicationBuffer that is not overwritten by the contents of publication.

Related concepts:

“Example 2: Managed MQ subscriber”

The managed MQ subscriber is the preferred pattern for most subscriber applications. The example requires *no* administrative definition of queues, topics or subscriptions.

“Example 3: Unmanaged MQ subscriber” on page 191

The unmanaged subscriber is an important class of subscriber application. With it, you combine the benefits of publish/subscribe with *control* of queuing and consumption of publications. The example demonstrates different ways of combining subscriptions and queues.

“Writing publisher applications” on page 173

Get started with writing publisher applications by studying two examples. The first is modeled as closely as possible on a point to point application putting messages on a queue, and the second demonstrates creating topics dynamically - a more common pattern for publisher applications.

Example 2: Managed MQ subscriber:

The managed MQ subscriber is the preferred pattern for most subscriber applications. The example requires *no* administrative definition of queues, topics or subscriptions.

This simplest kind of managed subscriber typically uses a *non-durable* subscription. The example focuses on a non-durable subscription. The subscription lasts only as long as the lifetime of the subscription handle from MQSUB. Any publications that match the topic string during the lifetime of the subscription are sent to the subscription queue (and possibly a retained publication if the flag MQSO_NEW_PUBLICATIONS_ONLY is not set or defaulted, an earlier publication matching the topic string was retained, and the publication was persistent or the queue manager has not terminated, since the publication was created).

You can also use a *durable* subscription with this pattern. Typically if a managed durable subscription is used it is done for *reliability* reasons, rather than to establish a subscription that, without any errors occurring, would outlive the subscriber. For more information about different life cycles associated with managed, unmanaged, durable and non-durable subscriptions see the related topics section.

Durable subscriptions are often associated with persistent publications, and non-durable subscriptions with non-persistent publications, but there is no necessary relationship between subscription durability and publication persistence. All four combinations of persistence and durability are possible.

For the managed non-durable case considered, the queue manager creates a subscription queue that is purged and deleted when the queue is closed. The publications are removed from the queue when the non-durable subscription is closed.

The valuable facets of the managed non-durable pattern exemplified by this code are as follows:

1. On demand subscription: the subscription topic string is dynamic. It is provided by the application when it runs.
2. Self managing queue: the subscription queue is self defining and managing.
3. Self managing subscription lifecycle: *non-durable* subscriptions only exist for the duration of the subscriber application.
 - If you define a *durable* managed subscription, then it results in a permanent subscription queue and publications continue to be stored on it with no subscriber programs being active. The queue manager deletes the queue (and clears any unretrieved publications from it) only after the application or administrator has chosen to delete the subscription. The subscription can be deleted using an administrative command, or by closing the subscription with the MQCO_REMOVE_SUB option.
 - Consider setting SubExpiry for durable subscriptions so that publications cease to be sent to the queue and the subscriber can consume any remaining publications before removing the subscription and causing the queue manager to delete the queue and any remaining publications on it.
4. Flexible topic string deployment: Subscription topic management is simplified by defining the root part of the subscription using an administratively defined topic. The root part of the topic tree is then hidden from the application. By hiding the root part an application can be deployed without the application inadvertently creating a topic tree that overlaps with another topic tree created by another instance, or another application.
5. Administered topics: by using a topic string in which the first part matches an administratively defined topic object, publications are managed according to the attributes of the topic object.
 - For example, if the first part of the topic string matches the topic string associated with a clustered topic object, then the subscription can receive publications from other members of the cluster
 - The selective matching of administratively defined topic objects and programmatically defined subscriptions enables you to combine the benefits of both. The administrator provides attributes for topics, and the programmer dynamically defines subtopics without being concerned about the management of topics.
 - It is the resultant topic string which is used to match the topic object that provides the attributes associated with the topic, and not necessarily the topic object named in sd.Objectname, although they typically turn out to be one and the same. See “Example 2: Publisher to a variable topic” on page 178.

By making the subscription durable in the example, publications continue to be sent to the subscription queue after the subscriber has closed the subscription with the MQCO_KEEP_SUB option . The queue continues to receive publications when the subscriber is not active. You can override this behavior by creating the subscription with the MQSO_PUBLICATIONS_ON_REQUEST option and using MQSUBRQ to request the retained publication.

The subscription can be resumed later by opening the subscription with the MQCO_RESUME option.

You can use the queue handle, Hobj, returned by MQSUB in a number of ways. The queue handle is used in the example to inquire on the name of the subscription queue. Managed queues are opened using the default model queues SYSTEM.NDURABLE.MODEL.QUEUE or SYSTEM.DURABLE.MODEL.QUEUE. You can override the defaults by providing your own durable and non-durable model queues on a topic by topic basis as properties of the topic object associated with the subscription.

Regardless of the attributes inherited from the model queues, you cannot reuse a managed queue handle to create an additional subscription. Nor can you obtain another handle for the managed queue by opening the managed queue a second time using the returned queue name. The queue behaves as if it has been opened for exclusive input .

Unmanaged queues are more flexible than managed queues. You can, for example share unmanaged queues, or define multiple subscriptions on the one queue. The next example, , demonstrates how to combine subscriptions with an unmanaged subscription queue.

Note: The compact coding style is intended for readability not production use.

The results are shown in Figure 26 on page 190.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>

void inquireQname(MQHCONN HConn, MQHOBJ Hobj, MQCHAR48 qName);

int main(int argc, char **argv)
{
    MQCHAR48 topicNameDefault    = "STOCKS";
    char    topicStringDefault[] = "IBM/PRICE";
    MQCHAR48 qmName = "";          /* Use default queue manager */
    MQCHAR48 qName = "";          /* Allocate to query queue name */
    char    publicationBuffer[101]; /* Allocate to receive messages */
    char    resTopicStrBuffer[151]; /* Allocate to resolve topic string */

    MQHCONN Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ Hobj = MQHO_NONE; /* publication queue handle */
    MQHOBJ Hsub = MQSO_NONE; /* subscription handle */
    MQLONG CompCode = MQCC_OK; /* completion code */
    MQLONG Reason = MQRC_NONE; /* reason code */
    MQLONG messlen = 0;
    MQSD sd = {MQSD_DEFAULT}; /* Subscription Descriptor */
    MQMD md = {MQMD_DEFAULT}; /* Message Descriptor */
    MQGMO gmo = {MQGMO_DEFAULT}; /* get message options */

    char * topicName = topicNameDefault;
    char * topicString = topicStringDefault;
    char * publication = publicationBuffer;
    char * resTopicStr = resTopicStrBuffer;
    memset(resTopicStr, 0, sizeof(resTopicStrBuffer));

    switch(argc){
        /* Replace defaults with args if provided */
        default:
            topicString = argv[2];
        case(2):
            if (strcmp(argv[1],"/") /* "/" invalid = No topic object */
                topicName = argv[1];
            else
                *topicName = '\0';
        case(1):
            printf("Optional parameters: topicName, topicString\nValues \"%s\" \"%s\"\n",
                topicName, topicString);
    }
}
```

Figure 24. Managed MQ subscriber - part 1: declarations and parameter handling.

There are some additional comments to make about the declarations in this example.

MQHOBJ Hobj = MQHO_NONE;

You cannot explicitly open a non-durable managed subscription queue to receive publications, but you do need to allocate storage for the object handle the queue manager returns when it opens the queue for you. It is important to initialize the handle to MQHO_OBJECT. This indicates to the queue manager that it needs to return a queue handle to the subscription queue.

MQSD sd = {MQSD_DEFAULT};

The new subscription descriptor, used in MQSUB.

MQCHAR48 qName;

Although the example doesn't require knowledge of the subscription queue, the example does inquire the name of the subscription queue - the MQINQ binding is a little awkward in the C language, so you might find this part of the example useful to study.

```

do {
    MQCONN(qmName, &Hconn, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    strncpy(sd.ObjectName, topicName, MQ_TOPIC_NAME_LENGTH);
    sd.ObjectString.VSPtr = topicString;
    sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;
    sd.Options = MQSO_CREATE | MQSO_MANAGED | MQSO_NON_DURABLE | MQSO_FAIL_IF QUIESCING ;
    sd.ResObjectString.VSPtr = resTopicStr;
    sd.ResObjectString.VSBufSize = sizeof(resTopicStrBuffer)-1;
    MQSUB(Hconn, &sd, &Hobj, &Hsub, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    gmo.Options = MQGMO_WAIT | MQGMO_NO_SYNCPOINT | MQGMO_CONVERT;
    gmo.WaitInterval = 10000;
    inquireQname(Hconn, Hobj, qName);
    printf("Waiting %d seconds for publications matching \"%s\" from \"%-0.48s\"\n",
        gmo.WaitInterval/1000, resTopicStr, qName);
    do {
        memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
        memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
        md.Encoding = MQENC_NATIVE;
        md.CodedCharSetId = MQCCSI_Q_MGR;
        memset(publicationBuffer, 0, sizeof(publicationBuffer));
        MQGET(Hconn, Hobj, &md, &gmo, sizeof(publicationBuffer)-1,
            publication, &messlen, &CompCode, &Reason);
        if (Reason == MQRC_NONE)
            printf("Received publication \"%s\"\n", publication);
    }
    while (CompCode == MQCC_OK);
    if (CompCode != MQCC_OK && Reason != MQRC_NO_MSG_AVAILABLE) break;
    MQCLOSE(Hconn, &Hsub, MQCO_REMOVE_SUB, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    MQDISC(&Hconn, &CompCode, &Reason);
} while (0);
printf("Completion code %d and Return code %d\n", CompCode, Reason);
return;
}
void inquireQname(MQHCONN Hconn, MQHOBJ Hobj, MQCHAR48 qName) {
#define _selectors 1
#define _intAttrs 1

    MQLONG select[_selectors] = {MQCA_Q_NAME}; /* Array of attribute selectors */
    MQLONG intAttrs[_intAttrs]; /* Array of integer attributes */
    MQLONG CompCode, Reason;
    MQINQ(Hconn, Hobj, _selectors, select, _intAttrs, intAttrs, MQ_Q_NAME_LENGTH, qName,
        &CompCode, &Reason);
    if (CompCode != MQCC_OK) {
        printf("MQINQ failed with Condition code %d and Reason %d\n", CompCode, Reason);
        strcpy(qName, "unknown queue");
    }
    return;
}

```

Figure 25. Managed MQ subscriber - part 2: code body.

```

W:\Subscribe2\Debug>solution2
Optional parameters: topicName, topicString
Values "STOCKS" "IBM/PRICE"
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from "SYSTEM.MANAGED.NDURABLE.48A0AC7403300020"
Received publication "150"
Completion code 0 and Return code 0

W:\Subscribe2\Debug>solution2 / NYSE/IBM/PRICE
Optional parameters: topicName, topicString
Values "" "NYSE/IBM/PRICE"
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from "SYSTEM.MANAGED.NDURABLE.48A0AC7403310020"
Received publication "150"
Completion code 0 and Return code 0

```

Figure 26. MQ subscriber

There are some additional comments to make about the code in this example.

strncpy(sd.ObjectName, topicName, MQ_Q_NAME_LENGTH);

If topicName is null or blank (*default value*), the topic name is not used to compute the resolved topic string.

sd.ObjectString.VSPtr = topicString;

Rather than solely use a predefined topic object, in this example the programmer provides a topic object and a topic string, that are combined by MQSUB. Notice the topic string is a MQCHARV structure.

sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;

An alternative to setting the length of a MQCHARV field.

sd.Options = MQSO_CREATE | MQSO_MANAGED | MQSO_NON_DURABLE | MQSO_FAIL_IF QUIESCING;

After defining the topic string, the sd.Options flags need the most careful attention. There are many options, the example specifies only the most commonly used ones. The other options use the default values.

1. As the subscription is *non-durable*, that is, it has a lifetime of the open subscription in the application, set the MQSO_CREATE flag. You can also set the (*default*) MQSO_NON_DURABLE flag for readability.
2. Complementing MQSO_CREATE is MQSO_RESUME. Both flags can be set together; the queue manager either creates a new subscription or resumes an existing subscription, whichever is appropriate. However, if you do specify MQSO_RESUME you must also initialize the MQCHARV structure for sd.SubName, even if there is no subscription to resume. Failure to initialize SubName results in a return code of 2440: MQRC_SUB_NAME_ERROR from MQSUB.

Note: MQSO_RESUME is always ignored for a non-durable managed subscription: but specifying it without initializing the MQCHARV structure for sd.SubName does cause the error.

3. In addition there is a third flag affecting how the subscription is opened, MQSO_ALTER. Given the correct permissions, the properties of a resumed subscription are changed to match other attributes specified in MQSUB.

Note: At least one of the MQSO_CREATE, MQSO_RESUME and MQSO_ALTER flags must be specified. See Options (MQLONG). There are examples of using all three flags in “Example 3: Unmanaged MQ subscriber” on page 191.

4. Set MQSO_MANAGED for the queue manager to manage the subscription for you automatically.

sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;

Optionally, omit setting the length of MQCHARV for null terminated strings and use the null terminator flag instead.

sd.ResObjectString.VSPtr = resTopicStr;

The resulting topic string is echoed in first printf in the program. Set up MQCHARV ResObjectString for IBM MQ to return the resolved string back to the program.

Note: `resTopicStringBuffer` is initialized to nulls in `memset(resTopicStr, 0, sizeof(resTopicStrBuffer))`. Returned topic strings do not end with a trailing null.

`sd.ResObjectString.VSBufSize = sizeof(resTopicStrBuffer)-1;`

Set the buffer size of the `sd.ResObjectString` to one less than its actual size. This prevents overwriting the null terminator that is provided, in case the resolved topic string fills the entire buffer.

Note: No error is returned if the topic string is longer than `sizeof(resTopicStrBuffer)-1`. Even if `VSLength > VSBufSiz` the length returned in `sd.ResObjectString.VSLength` is the length of the complete string and not necessarily the length of the returned string. Test `sd.ResObjectString.VSLength < sd.ResObjectString.VSBufSiz` to confirm the topic string is complete.

`MQSUB(Hconn, &sd, &Hobj, &Hsub, &CompCode, &Reason);`

The `MQSUB` function creates a subscription. If it is non-durable you are probably not interested in its name, though you can inspect its status in IBM MQ Explorer. You can provide the `sd.SubName` parameter as input, so you know what name to look for; you obviously have to avoid name clashes with other subscriptions.

`MQCLOSE(Hconn, &Hsub, MQCO_REMOVE_SUB, &CompCode, &Reason);`

Closing both the subscription and the subscription queue is optional. In the example the subscription is closed, but not the queue. The `MQCLOSE MQCO_REMOVE_SUB` option is the default in this case anyway as the subscription is non-durable. Using `MQCO_KEEP_SUB` is an error.

Note: the subscription *queue* is not closed by `MQSUB`, and its handle, `Hobj`, remains valid until the queue is closed by `MQCLOSE` or `MQDISC`. If the application terminates prematurely, the queue and subscription are cleaned up by the queue manager sometime after application termination.

Related concepts:

“Example 1: MQ Publication consumer” on page 182

The MQ Publication consumer is an IBM MQ message consumer that does not subscribe to topics itself.

“Example 3: Unmanaged MQ subscriber”

The unmanaged subscriber is an important class of subscriber application. With it, you combine the benefits of publish/subscribe with *control* of queuing and consumption of publications. The example demonstrates different ways of combining subscriptions and queues.

“Writing publisher applications” on page 173

Get started with writing publisher applications by studying two examples. The first is modeled as closely as possible on a point to point application putting messages on a queue, and the second demonstrates creating topics dynamically - a more common pattern for publisher applications.

Example 3: Unmanaged MQ subscriber:

The unmanaged subscriber is an important class of subscriber application. With it, you combine the benefits of publish/subscribe with *control* of queuing and consumption of publications. The example demonstrates different ways of combining subscriptions and queues.

The unmanaged pattern is more commonly associated with *durable* subscriptions than *non-durable*. Typically the lifecycle of a subscription created by an unmanaged subscriber is independent of the lifecycle of the subscribing application itself. By making the subscription durable the subscription receives publications even when no subscribing application is active.

You can create durable *managed* subscriptions to achieve the same result, but some applications require more flexibility and control over queues and messages than is possible with a managed subscription. For a durable managed subscription, the queue manager creates a permanent queue for the publications that match the subscription topic. It deletes the queue and associated publications when the subscription is deleted.

Typically durable *managed* subscriptions are used if the lifecycle of the application and the subscription is essentially the same, but hard to guarantee. By making the subscription durable, and having the publisher create persistent publications, there are no lost messages should the queue manager or subscriber terminate prematurely and need to be recovered.

The queue manager implicitly opens the durable managed subscription queue for a subscriber in such a way that shared processing of the queue is not possible. In addition, you cannot create more than one subscription for each managed queue and you may find the queues harder to manage because you have less control over the names of the queues. For these reasons, consider whether the *unmanaged* MQ subscriber is a better fit for applications requiring durable subscriptions than the *managed* MQ subscriber.

The code in Figure 29 on page 198 demonstrates an unmanaged durable subscription pattern. For illustration the code also creates unmanaged, non-durable subscriptions. This example illustrates the following pattern facets:

- On demand subscriptions: the subscription topic strings are dynamic. They are provided by the application when it runs.
- Simplified subscription topic management: subscription topic management is simplified by defining the root part of the subscription topic string using an administratively defined topic. This hides the root part of the topic tree from the application. By hiding the root part a subscriber can be deployed to different topic trees.
- Flexible subscription management: you can define a subscription either administratively, or create it on-demand in a subscriber program. There is no difference between administratively and programmatically created subscriptions, except an attribute that shows how the subscription was created. There is a third type of subscription that is created automatically by the queue manager for distribution of subscriptions. All subscriptions are displayed in the IBM MQ Explorer.
- Flexible association of subscriptions with queues: a predefined local queue is associated with a subscription by the MQSUB function. There are different ways to use MQSUB to associate subscriptions with queues:
 - Associate a subscription with a queue having *no* existing subscriptions, MQSO_CREATE + (Hobj from MQOPEN).
 - Associate a *new* subscription with a queue having existing subscriptions, MQSO_CREATE + (Hobj from MQOPEN).
 - Move an existing subscription to a different queue, MQSO_ALTER + (Hobj from MQOPEN).
 - Resume an existing subscription associated with an existing queue, MQSO_RESUME + (Hobj = MQHO_NONE), or MQSO_RESUME + (Hobj = from MQOPEN of queue with existing subscription).
 - By combining MQSO_CREATE | MQSO_RESUME | MQSO_ALTER in different combinations, you can cater for different input states of the subscription and the queue without having to code multiple versions of MQSUB with different sd.Options values.
 - Alternatively, by coding a specific choice of MQSO_CREATE | MQSO_RESUME | MQSO_ALTER the queue manager returns an error (Table 25 on page 194) if the states of the subscription and queue provided as input to MQSUB are inconsistent with the value of sd.Options. Figure 35 on page 201 shows the results of issuing MQSUB for Subscription X with different individual settings of the sd.Options flag, and passing it three different object handles.

Explore different inputs to the example program in Figure 28 on page 196 to become familiar with these different kinds of errors. One common error, RC = 2440, that is not included in the cases listed in the table, is a subscription name error. It is commonly caused by passing a null or invalid subscription name with MQSO_RESUME or MQSO_ALTER.

- Multiprocessing: You can share among many consumers the work of reading publications. The publications all go onto the single queue associated with the subscription topic. Consumers have a choice of opening the queue directly using MQOPEN or resuming the subscription using MQSUB.

- Subscription concentration: multiple subscriptions can be created on the same queue. Be cautious with this capability as it can lead to overlapping subscriptions, and receiving the same publication multiple times. The MQSO_GROUP_SUB option eliminates duplicate publications caused by overlapping subscriptions.
- Subscriber and consumer separation: As well as the three consumer models illustrated in the examples, another model is to separate the consumer from the subscriber. It is a variation of the unmanaged MQ Subscriber, but rather than issue the MQOPEN and MQSUB in the same program, one program subscribes to publications, and another program consumes them. For example, the subscriber might be part of a publish/subscribe cluster and the consumer attached to a queue manager outside the queue manager cluster. The consumer receives publications through standard distributed queuing by defining the subscription queue as a remote queue definition.

Understanding the behavior of MQSO_CREATE | MQSO_RESUME | MQSO_ALTER is important, especially if you plan to simplify your code by using combinations of these options. Study the table Table 25 on page 194 that shows the results of passing different queue handles to MQSUB, and the results of running the example program shown in Figure 30 on page 199 to Figure 35 on page 201.

The scenario used to construct the table has one subscription X and two queues, A and B. The subscription name parameter sd.SubName is set to X, the name of a subscription attached to queue A. Queue B has no subscription attached to it.

In Table 25 on page 194, MQSUB is passed subscription X and the queue handle to queue A. The results from subscription options are as follows:

- MQSO_CREATE fails because the queue handle corresponds to the queue A which already has a subscription to X. Contrast this behavior to the successful call. That call succeeds because queue B does not have a subscription to X attached to it.
- MQSO_RESUME succeeds because the queue handle corresponds to the queue A which already has a subscription to X. In contrast, the call fails where the subscription X does not exist on queue A.
- MQSO_ALTER behaves in a similar way to MQSO_RESUME with respect to opening the subscription and queue. However if the attributes contained within the subscription descriptor passed to MQSUB differ from the attributes of the subscription, MQSO_RESUME fails, whereas MQSO_ALTER succeeds as long as the program instance has permission to alter the attributes. Note that you can never change the topic string in a subscription; but rather than return an error, MQSUB ignores the topic name and topic string values in the subscription descriptor and uses the values in the existing subscription.

Next, look at Table 25 on page 194 where MQSUB is passed subscription X and the queue handle to queue B. The results from subscription options are as follows:

- MQSO_CREATE succeeds and creates subscription X on queue B because this is a new subscription on queue B.
- MQSO_RESUME fails. MQSUB looks for subscription X on queue B and does not find it, but rather than returning *RC = 2428 - subscription X does not exist*, it returns *RC = 2019 - Subscription queue does not match queue object handle*. The behavior of the third option MQSO_ALTER suggests the reason for this unexpected error. MQSUB expects the queue handle to point to a queue with a subscription. It checks this first before checking whether the subscription named in sd.SubName exists.
- MQSO_ALTER succeeds, and moves the subscription from queue A to queue B.

A case that is not shown in the table is if the subscription name of the subscription on queue A does not match the subscription name in sd.SubName. That call fails with a *RC = 2428 - subscription X does not exist on Queue A*.

Table 25. Errors from MQSUB with different queue handles and subscription combinations

	<p>Queue A Subscription X</p> <p>Queue B No subscription</p>	<p>Queue A No subscription</p> <p>Queue B No subscription</p>
Hobj for Queue A passed to MQSUB	<p>MQSO_CREATE RC = 2432 - Subscription X already exists on Queue A</p> <p>MQSO_RESUME Resumes subscription X on Queue A</p> <p>MQSO_ALTER Resumes subscription X on Queue A and makes permitted alterations</p>	<p>MQSO_CREATE Creates subscription X on Queue A</p> <p>MQSO_RESUME RC = 2428 - Subscription X does not exist on Queue A</p> <p>MQSO_ALTER RC = 2428 - Subscription X does not exist on Queue A</p>
Hobj for Queue B passed to MQSUB	<p>MQSO_CREATE Creates new subscription X on Queue B</p> <p>MQSO_RESUME RC = 2019 - Subscription queue does not match queue object handle</p> <p>MQSO_ALTER Move subscription X from Queue A to Queue B</p>	<p>MQSO_CREATE Creates new subscription X on Queue B</p> <p>MQSO_RESUME RC = 2428 - subscription X does not exist on Queue B</p> <p>MQSO_ALTER RC = 2428 - subscription X does not exist on Queue B</p>
MQHO_NONE passed to MQSUB	<p>MQSO_CREATE RC = 2019 - Bad object handle: set MQSO_MANAGED flag to create a managed subscription and create a managed queue</p> <p>MQSO_RESUME Resumes subscription X on Queue A and returns Hobj to Queue A</p> <p>MQSO_ALTER Resumes subscription X on Queue A, returns Hobj to Queue A and makes permitted alterations</p>	<p>MQSO_CREATE RC = 2019 - Bad object handle: set MQSO_MANAGED flag to create a managed subscription and create a managed queue</p> <p>MQSO_RESUME RC = 2428 - No subscription X</p> <p>MQSO_ALTER RC = 2019 - Bad object handle: No queue A or B</p>

Note: The compact coding style is intended for readability not production use.

```

#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>

void inquireQname(MQHCONN HConn, MQHOBJ Hobj, MQCHAR48 qName);

int main(int argc, char **argv)
{
    MQCHAR48 topicNameDefault          = "STOCKS";
    char      topicStringDefault[]      = "IBM/PRICE";
    char      subscriptionNameDefault[] = "IBMSTOCKPRICESUB";
    char      subscriptionQueueDefault[] = "STOCKTICKER";
    char      publicationBuffer[101];   /* Allocate to receive messages */
    char      resTopicStrBuffer[151];   /* Allocate to resolve topic string */
    MQCHAR48 qmName = "";               /* Default queue manager */
    MQCHAR48 qName = "";               /* Allocate storage for MQINQ */

    MQHCONN  Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ   Hobj = MQHO_NONE;           /* subscription queue handle */
    MQHOBJ   Hsub = MQSO_NONE;          /* subscription handle */
    MQLONG   CompCode = MQCC_OK;        /* completion code */
    MQLONG   Reason = MQRC_NONE;        /* reason code */
    MQLONG   messlen = 0;
    MQOD     od = {MQOD_DEFAULT};       /* Unmanaged subscription queue */
    MQSD     sd = {MQSD_DEFAULT};       /* Subscription Descriptor */
    MQMD     md = {MQMD_DEFAULT};       /* Message Descriptor */
    MQGMO    gmo = {MQGMO_DEFAULT};     /* get message options */
    MQLONG   sdOptions = MQSO_CREATE | MQSO_RESUME | MQSO_DURABLE | MQSO_FAIL_IF QUIESCING;

    char *   topicName = topicNameDefault;
    char *   topicString = topicStringDefault;
    char *   subscriptionName = subscriptionNameDefault;
    char *   subscriptionQueue = subscriptionQueueDefault;
    char *   publication = publicationBuffer;
    char *   resTopicStr = resTopicStrBuffer;
    memset(resTopicStrBuffer, 0, sizeof(resTopicStrBuffer));

```

Figure 27. Unmanaged MQ subscriber - part 1: declarations.

```

    switch(argc){
        /* Replace defaults with args if provided */
    default:
        switch((argv[5][0])) {
    case('A'): sdOptions = MQSO_ALTER | MQSO_DURABLE | MQSO_FAIL_IF QUIESCING;
                break;
    case('C'): sdOptions = MQSO_CREATE | MQSO_DURABLE | MQSO_FAIL_IF QUIESCING;
                break;
    case('R'): sdOptions = MQSO_RESUME | MQSO_DURABLE | MQSO_FAIL_IF QUIESCING;
                break;
    default:    ;
        }
    case(5):
        if (strcmp(argv[4],"/") /* "/" invalid = No subscription */
            subscriptionQueue = argv[4];
        else {
            *subscriptionQueue = '\0';
            if (argc > 5) {
                if (argv[5][0] == 'C') {
                    sdOptions = sdOptions + MQSO_MANAGED;
                }
            }
            else
                sdOptions = sdOptions + MQSO_MANAGED;
        }
    case(4):
        if (strcmp(argv[3],"/") /* "/" invalid = No subscription */
            subscriptionName = argv[3];
        else {
            *subscriptionName = '\0';
            sdOptions = sdOptions - MQSO_DURABLE;
        }
    case(3):
        if (strcmp(argv[2],"/") /* "/" invalid = No topic string */
            topicString = argv[2];
        else
            *topicString = '\0';
    case(2):
        if (strcmp(argv[1],"/") /* "/" invalid = No topic object */
            topicName = argv[1];
        else
            *topicName = '\0';
    case(1):
        sd.Options = sdOptions;
        printf("Optional parameters: "
            printf("topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|C(reate)|R(esume)\n");
        printf("Values \"%-.48s\" \"%s\" \"%s\" \"%-.48s\" sd.Options=%d\n",
            topicName, topicString, subscriptionName, subscriptionQueue, sd.Options);
    }
}

```

Figure 28. Unmanaged MQ subscriber - part 2: parameter handling.

Additional comments about the parameter handling in this example are as follows:

switch((argv[5][0]))

You have the choice of entering A lter | C reate | R esume in parameter 5, to test the effect of overriding part of the MQSUB option setting used by default in the example. The default setting used by the example is MQSO_CREATE | MQSO_RESUME | MQSO_DURABLE.

Note: Setting MQSO_ALTER or MQSO_RESUME without setting MQSO_DURABLE is an error, and sd.SubName must be set and refer to a subscription that can be resumed or altered.

***subscriptionQueue = '\0';**

sdOptions = sdOptions + MQSO_MANAGED;

If the default subscription queue, STOCKTICKER is replaced by a null string then as long as MQSO_CREATE is set, the example sets the MQSO_MANAGED flag and creates a dynamic subscription queue. If Alter or Resume are set in the fifth parameter the behavior of the example will depend on the value of subscriptionName.

***subscriptionName = '\0';**

sdOptions = sdOptions - MQSO_DURABLE;

If the default subscription, IBMSTOCKPRICESUB, is replaced by a null string then the example removes the MQSO_DURABLE flag. If you run the example providing the default values for the other parameters an additional temporary subscription destined to STOCKTICKER is created and receives duplicate publications. Next time you run the example, without any parameters, you receive just one publication again.

```

do {
    MQCONN(qmName, &Hconn, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    if (strlen(subscriptionQueue)) {
        strncpy(od.ObjectName, subscriptionQueue, MQ_Q_NAME_LENGTH);
        MQOPEN(Hconn, &od, MQOO_INPUT_AS_Q_DEF | MQOO_FAIL_IF QUIESCING | MQOO_INQUIRE,
            &Hobj, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
    }
    strncpy(sd.ObjectName, topicName, MQ_TOPIC_NAME_LENGTH);
    sd.ObjectString.VSPtr = topicString;
    sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;
    sd.SubName.VSPtr = subscriptionName;
    sd.SubName.VSLength = MQVS_NULL_TERMINATED;
    sd.ResObjectString.VSPtr = resTopicStr;
    sd.ResObjectString.VSBufSize = sizeof(resTopicStrBuffer)-1;
    MQSUB(Hconn, &sd, &Hobj, &Hsub, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    gmo.Options = MQGMO_WAIT | MQGMO_NO_SYNCPOINT | MQGMO_CONVERT;
    gmo.WaitInterval = 10000;
    inquireQname(Hconn, Hobj, qName);
    printf("Waiting %d seconds for publications matching \"%s\" from %-0.48s\n",
        gmo.WaitInterval/1000, resTopicStr, qName);
    do {
        memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
        memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
        md.Encoding = MQENC_NATIVE;
        md.CodedCharSetId = MQCCSI_Q_MGR;
        MQGET(Hconn, Hobj, &md, &gmo, sizeof(publication), publication, &messlen, &CompCode, &Reason);
        if (Reason == MQRC_NONE)
            printf("Received publication \"%s\"\n", publication);
    }
    while (CompCode == MQCC_OK);
    if (CompCode != MQCC_OK && Reason != MQRC_NO_MSG_AVAILABLE) break;
    MQCLOSE(Hconn, &Hsub, MQCO_NONE, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    MQDISC(&Hconn, &CompCode, &Reason);
} while (0);
printf("Completion code %d and Return code %d\n", CompCode, Reason);
}
void inquireQname(MQHCONN Hconn, MQHOBJ Hobj, MQCHAR48 qName) {
#define _selectors 1
#define _intAttrs 1

    MQLONG select[_selectors] = {MQCA_Q_NAME}; /* Array of attribute selectors */
    MQLONG intAttrs[_intAttrs]; /* Array of integer attributes */
    MQLONG CompCode, Reason;
    MQINQ(Hconn, Hobj, _selectors, select, _intAttrs, intAttrs, MQ_Q_NAME_LENGTH, qName, &CompCode, &Reason);
    if (CompCode != MQCC_OK) {
        printf("MQINQ failed with Condition code %d and Reason %d\n", CompCode, Reason);
        strncpy(qName, "unknown queue", MQ_Q_NAME_LENGTH);
    }
    return;
}

```

Figure 29. Unmanaged MQ subscriber - part 3: code body.

Additional comments on the code in this example are as follows:

if (strlen(subscriptionQueue))

If there is no subscription queue name then the example uses MQHO_NONE as the value of Hobj.

MQOPEN(...);

The subscription queue is opened and the queue handle saved in Hobj.

MQSUB(Hconn, &sd, &Hobj, &Hsub, &CompCode, &Reason);

The subscription is opened using the Hobj passed from MQOPEN (or MQHO_NONE if there is no subscription queue name). An unmanaged queue can be resumed without explicitly opening it with an MQOPEN.

MQCLOSE(Hconn, &Hsub, MQCO_NONE, &CompCode, &Reason);

The subscription is closed using the subscription handle. Depending on whether the subscription is durable or not, the subscription is closed with an implicit MQCO_KEEP_SUB or MQCO_REMOVE_SUB. You can close a durable subscription with MQCO_REMOVE_SUB, but you cannot close a non-durable subscription with MQCO_KEEP_SUB. The action of MQCO_REMOVE_SUB is to remove the subscription which stops any further publications being sent to the subscription queue.

MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);

No special action is taken if the subscription is unmanaged. If the queue is managed and the subscription closed with either an explicit or implicit MQCO_REMOVE_SUB, then all publications are purged from the queue and queue deleted at this point.

Results from the example illustrate aspects of publish/subscribe:

In Figure 30 the example starts by publishing 130 on the NYSE/IBM/PRICE topic.

```
W:\Subscribe3\Debug>..\Publish2\Debug\publishstock
Provide parameters: TopicObject TopicString Publication
Publish "130" to topic "STOCKS" and topic string "IBM/PRICE"
Published "130" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0
```

Figure 30. Publish 130 to NYSE/IBM/PRICE

In Figure 31 execution of the example using default parameters receives the retained publication 130. The provided topic object and topic string are ignored, as shown in Figure 35 on page 201. The topic object and topic string are always taken from the subscription object, when one is provided, and the topic string is immutable. The actual behavior of the example depends on the choice or combination of MQSO_CREATE, MQSO_RESUME, and MQSO_ALTER. In this example MQSO_RESUME is the option selected.

```
W:\Subscribe3\Debug>solution3
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|C(reate)|R(esume)
Values "STOCKS" "IBM/PRICE" "IBMSTOCKPRICESUB" "STOCKTICKER" sd.Options=8206
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from STOCKTICKER
Received publication "130"
Completion code 0 and Return code 0
```

Figure 31. Receive the retained publication

In (Figure 32 on page 200) no publications are received, because the durable subscription has already received the retained publication. In this example, the subscription is resumed by providing only the subscription name without the queue name. If the queue name was provided, the queue would be opened first and the handle passed to MQSUB.

Note: The 2038 error from MQINQ is due to the implicit MQOPEN of STOCKTICKER by MQSUB not including the MQOO_INQUIRE option. Avoid the 2038 return code from MQINQ by opening the queue explicitly.

```

W:\Subscribe3\Debug>solution3 STOCKS IBM/PRICE IBMSTOCKPRICESUB / Resume
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|C(reate)|R(esume)
Values "STOCKS" "IBM/PRICE" "IBMSTOCKPRICESUB" "" sd.Options=8204
MQINQ failed with Condition code 2 and Reason 2038
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from unknown queue
Completion code 0 and Return code 0

```

Figure 32. Resume subscription

In Figure 33, the example creates a non-durable unmanaged subscription using STOCKTICKER as the destination. Because this is a new subscription, it receives the retained publication.

```

W:\Subscribe3\Debug>solution3 STOCKS IBM/PRICE / STOCKTICKER Create
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|C(reate)|R(esume)
Values "STOCKS" "IBM/PRICE" "" "STOCKTICKER" sd.Options=8194
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from STOCKTICKER
Received publication "130"
Completion code 0 and Return code 0

```

Figure 33. Receive retained publication with new unmanaged non durable subscription

In Figure 34, to demonstrate overlapping subscriptions, another publication is sent, changing the retained publication. Next, a new non-durable, unmanaged subscription is created by not providing a subscription name. The retained publication is received twice, once for the new subscription, and once for the durable IBMSTOCKPRICESUB subscription that is still active on the STOCKTICKER queue. The example is an illustration that it is the queue that has subscriptions, and not the application. Despite not referring to the IBMSTOCKPRICESUB subscription in this invocation of the application, the application receives the publication twice: once from the durable subscription that was created administratively, and once from the non-durable subscription created by the application itself.

```

W:\Subscribe3\Debug>..\..\Publish2\Debug\publishstock
Provide parameters: TopicObject TopicString Publication
Publish "130" to topic "STOCKS" and topic string "IBM/PRICE"
Published "130" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

W:\Subscribe3\Debug>solution3 STOCKS IBM/PRICE / STOCKTICKER Create
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|C(reate)|R(esume)
Values "STOCKS" "IBM/PRICE" "" "STOCKTICKER" sd.Options=8194
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from STOCKTICKER
Received publication "130"
Received publication "130"
Completion code 0 and Return code 0

```

Figure 34. Overlapping subscriptions

In Figure 35 on page 201 the example demonstrates that providing a new topic string and an existing subscription does not result in a changed subscription.

1. In the first case, Resume resumes the existing subscription, as you might expect, and ignores the changed topic string.
2. In the second case, Alter causes an error, RC = 2510, Topic not alterable.
3. In the third example, Create causes an error RC = 2432, Sub already exists.


```

W:\Subscribe3\Debug>solution3 "" NASDAQ/IBM/PRICE IBMSTOCKPRICESUB STOCKTICKER Resume
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|C(reate)|R(esume)
Values "" "NASDAQ/IBM/PRICE" "IBMSTOCKPRICESUB" "STOCKTICKER" sd.Options=8204
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from STOCKTICKER
Received publication "130"
Completion code 0 and Return code 0

W:\Subscribe3\Debug>solution3 "" NASDAQ/IBM/PRICE IBMSTOCKPRICESUB STOCKTICKER Alter
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|C(reate)|R(esume)
Values "" "NASDAQ/IBM/PRICE" "IBMSTOCKPRICESUB" "STOCKTICKER" sd.Options=8201
Completion code 2 and Return code 2510

W:\Subscribe3\Debug>solution3 "" NASDAQ/IBM/PRICE IBMSTOCKPRICESUB STOCKTICKER Create
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|C(reate)|R(esume)
Values "" "NASDAQ/IBM/PRICE" "IBMSTOCKPRICESUB" "STOCKTICKER" sd.Options=8202
Completion code 2 and Return code 2432

```

Figure 35. Subscription topics cannot be changed

Related concepts:

“Example 1: MQ Publication consumer” on page 182

The MQ Publication consumer is an IBM MQ message consumer that does not subscribe to topics itself.

“Example 2: Managed MQ subscriber” on page 185

The managed MQ subscriber is the preferred pattern for most subscriber applications. The example requires *no* administrative definition of queues, topics or subscriptions.

“Writing publisher applications” on page 173

Get started with writing publisher applications by studying two examples. The first is modeled as closely as possible on a point to point application putting messages on a queue, and the second demonstrates creating topics dynamically - a more common pattern for publisher applications.

Publish/subscribe lifecycles:

Consider the lifecycles of topics, subscriptions, subscribers, publications, publishers and queues in designing publish/subscribe applications.

The lifecycle of an object, such as a subscription, starts with its creation and ends with its deletion. It may also include other states and changes that it goes through, such as temporary suspension, having parent and children topics, expiration and deletion.

Traditionally IBM MQ objects such as queues are created administratively, or by administrative programs using Programmable Command Format (PCF). Publish/subscribe is different in providing the MQSUB and MQCLOSE API verbs to create and delete subscriptions, having the concept of managed subscriptions that not only create and delete queues, but also clean up unconsumed messages, and having associations between administratively created topic objects and programmatically or administratively created topic strings.

This functional richness caters for a wide range of publish/subscribe requirements, and also simplifies designing some common patterns of publish/subscribe application. Managed subscriptions, for example, simplify both the programming and administration of a subscription that is intended to last only as long as the program that created it. Unmanaged subscriptions simplify programming where there is a looser connection between subscribing and consuming publications. Centrally created subscriptions are useful where the pattern is one of routing publication traffic to consumers based on a centralized model of control, for example sending flight information to automated gates, whereas programmatically created subscriptions might be used if gate staff are responsible for subscribing to the passengers records for that flight, by entering a flight number at a gate.

In this last example a managed durable subscription might be appropriate: managed, because the subscriptions are being created very often, and have a clear endpoint when the gate closes and the subscription can be programmatically removed; durable, to avoid losing a passenger record due to the

gate subscriber program going down for one reason or another². To initiate the publication of passenger records to the gate, a possible design would be for the gate application to both subscribe to the passenger records using the gate number, and publish the gate opening event using the gate number. The publisher responds to the gate opening event by publishing the passenger records - which might then also go to other interested parties, such as billing, to record the flight is taking place, and to customer services, to text notifications to passengers' mobile phones of the gate number.

The centrally managed subscription might use a durable unmanaged model, routing passenger lists to the gate using a predefined queue for each gate.

The following three examples of publish/subscribe lifecycles illustrate how managed non-durable, managed durable, and unmanaged durable subscribers interact with subscriptions, topics, queues, publishers and the queue manager, and how the responsibilities might be divided between administration and the subscriber programs.

Managed non-durable subscriber

Figure 36 on page 203 shows an application creating a managed non-durable subscription, getting two messages that are published to the topic identified in the subscription, and terminating. The interactions labeled in an italic gray font with dotted arrows are implicit.

There are some points to note.

1. The application creates a subscription on a topic that has already been published to twice. When the subscriber receives its first publication, it receives the *second* publication which is the currently retained publication.
2. The queue manager creates a temporary subscription queue as well as creating a subscription for the topic.
3. The subscription has an expiry. When the subscription expires no more publications on the topic are sent to this subscription, but the subscriber continues to get messages published before the subscription expired. Publication expiry is not affected by subscription expiry.
4. The fourth publication is not placed on the subscription queue and consequently the last MQGET does not return a publication.
5. Although the subscriber closes its subscription, it does not close its connection to the queue or the queue manager.
6. The queue manager cleans up shortly after the application terminates. Because the subscription is managed and non-durable, the subscription queue is deleted.

2. The publisher must send the passenger records as persistent messages to avoid other possible failures, of course.

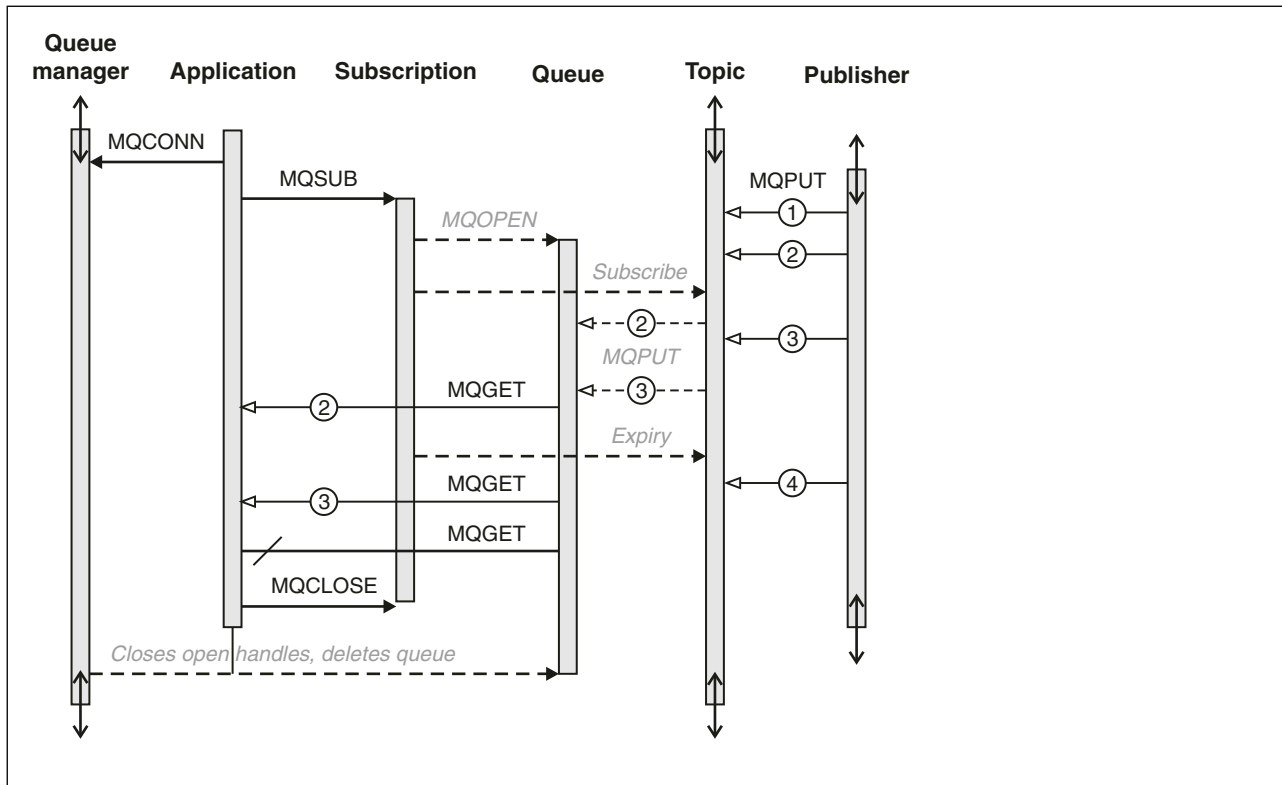


Figure 36. Managed non-durable subscriber lifelines

Managed durable subscriber

The managed durable subscriber takes the previous example a step further, and shows a managed subscription surviving the termination and restart of the subscribing application.

There are some new points to note.

1. In this example, unlike the last, the publication topic did not exist before it was defined in the subscription.
2. The first time the subscriber terminates, it closes the subscription with the option MQCO_KEEP_SUB. That is the default behavior for implicitly closing a managed durable subscription.
3. When the subscriber resumes the subscription, the subscription queue is reopened.
4. The new publication 2, placed on the queue before it is reopened, is available to MQGET, even after the subscription has been removed.

Even though the subscription is durable, the subscriber reliably receives all messages sent by the publisher only if *both* the subscription is durable and the messages persistent. Message persistence depends on the setting of the Persistent field in the MQMD of the message sent by the publisher. A subscriber has no control over this.

5. Closing the subscription with the flag MQCO_REMOVE_SUB removes the subscription, stopping any further publications being placed on the subscription queue. When the subscription queue is closed, then the queue manager removes the unread publication 3, and then deletes the queue. The action is equivalent to administratively deleting the subscription.

Note: Do not delete the queue manually, or issue MQCLOSE with the option MQCO_DELETE, or MQCO_PURGE_DELETE. The visible implementation details of a managed subscription is not part of the supported IBM MQ interface. The queue manager manage cannot manage a subscription reliably unless it has complete control.




Figure 37. Managed durable subscriber lifelines

Unmanaged durable subscriber

An administrator is added in the third example: the unmanaged durable subscriber. It is a good example to show how the administrator might interact with a publish/subscribe application.

The points to note are listed.

1. The publisher puts a message, 1, to a topic that later becomes associated with the topic object that is used for subscription. The topic object defines a topic string that matches the topic that was published to by using wildcards.
2. The topic has a retained publication.
3. The administrator creates a topic object, a queue and a subscription. The topic object and queue need to be defined before the subscription.
4. The application opens the queue associated with the subscription and passes MQSUB the handle of the queue. It could, alternatively, simply open the subscription, passing it the queue handle MQHO_NONE. The converse is not true, it cannot resume a subscription by passing it only queue handle without a subscription name - a queue might have multiple subscriptions.
5. The application opens the subscription using the option MQSO_RESUME even though it is the first time it has opened the subscription. It is resuming an administratively created subscription.

- The subscriber receives the retained publication, 1. Publication 2, although published before any publications were received by the subscriber, was published after the subscription started, and is the second publication on the subscription queue.

Note: If the retained publication is not published as a persistent message, then it is lost after queue manager restart.

- In this example the subscription is durable. It is possible for a program to create an unmanaged non-durable subscription; it should be obvious this is not something an administrator can do.
- The effect of the option MQCO_REMOVE_SUB on closing the subscription is to remove the subscription just as if the administrator had deleted it. This stops any further publications being sent to the queue, but does not affect publications that are already on the queue, even when the queue is closed, unlike a *managed* durable subscription.
- The administrator later deletes the remaining message, 3, and deletes the queue.

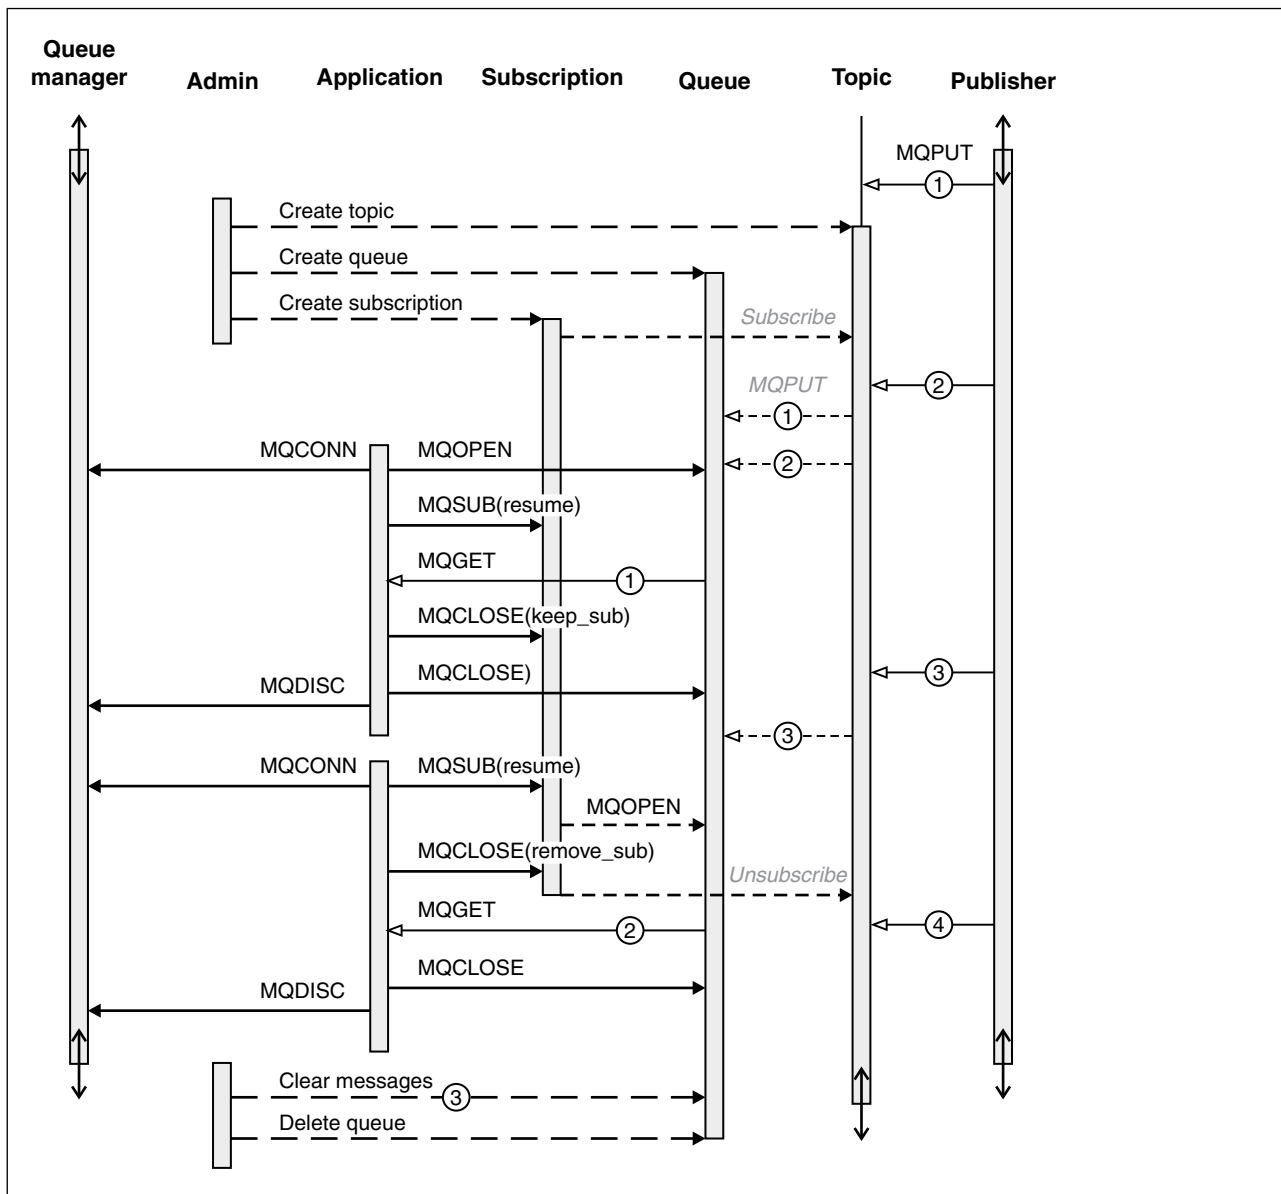


Figure 38. Unmanaged durable subscriber lifelines

A normal pattern for an unmanaged subscription is for queue and subscription housekeeping to be performed by the administrator. Typically one would not attempt to emulate the behavior of a managed subscriber and tidy up queues and subscriptions programmatically in application code. If you find yourself needing to write management logic, question whether you can achieve the same results using a managed pattern. It is not easy to write tightly synchronized, completely reliable management code. It is easier to tidy up later, either manually, or using an automated management program, when you can be sure that messages, subscriptions, and queues can be simply deleted, regardless of their state.

Publish/subscribe message properties:

Several message properties relate to IBM MQ publish/subscribe messaging.

PubAccountingToken

This is the value that will be in the AccountingToken field of the Message Descriptor (MQMD) of all publication messages matching this subscription. AccountingToken is part of the identity context of the message. For more information about message context, see “Message context” on page 40. For more information about the AccountingToken field in the MQMD, see AccountingToken.

PubApplIdentityData

This is the value that will be in the ApplIdentityData field of the Message Descriptor (MQMD) of all publication messages matching this subscription. ApplIdentityData is part of the identity context of the message. For more information about message context, see “Message context” on page 40. For more information about the ApplIdentityData field in the MQMD, see ApplIdentityData.

If the option MQSO_SET_IDENTITY_CONTEXT is not specified, the ApplIdentityData which will be set in each message published for this subscription is blanks, as default context information.

If the option MQSO_SET_IDENTITY_CONTEXT is specified, the PubApplIdentityData is being generated by the user and this field is an input field which contains the ApplIdentityData to be set in each publication for this subscription.

PubPriority

This is the value that will be in the Priority field of the Message Descriptor (MQMD) of all publication messages matching this subscription. For more information about the Priority field in the MQMD, see Priority.

The value must be greater than or equal to zero; zero is the lowest priority. The following special values can also be used:

- MQPRI_PRIORITY_AS_Q_DEF - When a subscription queue is provided in the Hobj field in the MQSUB call, and is not a managed handle, then the priority for the message is taken from the DefPriority attribute of this queue. If the queue so identified is a cluster queue, or there is more than one definition in the queue-name resolution path, the priority is determined when the publication message is put to the queue as described for Priority in the MQMD. If the MQSUB call uses a managed handle, the priority for the message is taken from the DefPriority attribute of the model queue associated with the topic subscribed to.
- MQPRI_PRIORITY_AS_PUBLISHED - The priority for the message is the priority of the original publication. This is the initial value of this field.

SubCorrelId

Attention: a correlation identifier can only be passed between queue managers in a publish/subscribe cluster, not a hierarchy.

All publications sent to match this subscription will contain this correlation identifier in the message descriptor. If multiple subscriptions use the same queue to get their publications from, using MQGET by correlation ID allows only publications for a specific subscription to be obtained. This correlation identifier can either be generated by the queue manager or by the user.

If the option MQSO_SET_CORREL_ID is not specified, the correlation identifier is generated by the queue manager and this field is an output field which contains the correlation identifier which will be set in each message published for this subscription.

If the option MQSO_SET_CORREL_ID is specified, the correlation identifier is being generated by the user and this field is an input field which contains the correlation identifier to be set in each publication for this subscription. In this case, if the field contains MQCL_NONE, the correlation identifier which will be set in each message published for this subscription will be the correlation identifier created by the original put of the message.

If the option MQSO_GROUP_SUB is specified and the correlation identifier specified is the same as an existing grouped subscription using the same queue and an overlapping topic string, only the most significant subscription in the group is provided with a copy of the publication.

SubUserData

This is the subscription user data. The data provided on the subscription in this field will be included as the MQSubUserData message property of every publication sent to this subscription.

Publication properties

Table 26 lists the publication properties that are provided with a publication message.

You can access these properties directly from the **MQRFH2** folder, or retrieve them using MQINQMP. MQINQMP accepts either the property name or **MQRFH2** name as the name of the property to inquire on.

Table 26. Publication properties

Property name	MQRFH2 name	Type	Description
MQTopicString	mmps.Top	MQTYPE_STRING	Topic string
MQSubUserData	mmps.Sud	MQTYPE_STRING	Subscriber user data
MQIsRetained	mmps.Ret	MQTYPE_BOOLEAN	Retained publication
MQPubOptions	mmps.Pub	MQTYPE_INT32	Publication options
MQPubLevel	mmps.Pb1	MQTYPE_INT32	Publication level
MQPubTime	mmpse.Pts	MQTYPE_STRING	Publication time
MQPubSeqNum	mmpse.Seq	MQTYPE_INT32	Publication sequence number
MQPubStrIntData	mmpse.Sid	MQTYPE_STRING	String/Integer data added by the publisher
MQPubFormat	mmpse.Pfmt	MQTYPE_INT32	Message format: MQRFH1 MQRFH2 PCF

Message ordering:

For a particular topic, messages are published by the queue manager in the same order as they are received from publishing applications (subject to reordering based on message priority).

Message ordering normally means that each subscriber receives messages from a particular queue manager, on a particular topic, from a particular publisher in the order that they are published by that publisher.

However, as with all IBM MQ messages, it is possible for messages, occasionally, to be delivered out of order. This can happen in the following situations:

- If a link in the network goes down and subsequent messages are rerouted along another link
- If a queue becomes temporarily full, or put-inhibited, so that a message is put to a dead-letter queue and therefore delayed, while subsequent messages pass straight through.
- If the administrator deletes a queue manager when publishers and subscribers are still operating, causing queued messages to be put to the dead-letter queue and subscriptions to be interrupted.

If these circumstances cannot occur, publications are always delivered in order.

Note: It is not possible to use grouped or segmented messages with Publish/Subscribe.

Intercepting publications:

You can intercept a publication, modify it, and then republish it before it reaches any other subscriber.

You might want to intercept a publication before it reaches a subscriber in order to do one of the following actions:

- Attach additional information to the message
- Block the message
- Transform the message

You can perform the same operation on each message or vary the operation depending on the subscription, the message, or the message header.

Related information:

MQ_PUBLISH_EXIT - Publish exit

Subscription levels:

Set the subscription level of a subscription to intercept a publication before it reaches its final subscribers. An intercepting subscriber subscribes at a higher subscription level, and republishes at a lower publication level. Build a chain of intercepting subscribers to perform message processing on a publication before it is delivered to final subscribers.




Figure 39. Sequence of intercepting subscribers

To intercept a publication, use the **MQSD** `SubLevel` attribute. After a message has been intercepted, it can be transformed and then republished at a lower publication level by changing the **MQPMO** `PubLevel` attribute. The message then goes to the final subscribers, or it is intercepted again by an intermediate subscriber at a lower subscription level.

The intercepting subscriber typically transforms a message before republishing it. A sequence of intercepting subscribers forms a message flow. Alternatively, you might not republish the intercepted publication: Subscribers at lower subscription levels would not receive the message.

Ensure that the interceptor receives publications before any other subscribers. Set the subscription level of the interceptor higher than other subscribers. By default, subscribers have a `SubLevel` of 1. The highest value is 9. A publication must start with a `PubLevel` at least as high as the highest `SubLevel`. Publish initially with the default `PubLevel` of 9.

- If you have one intercepting subscriber on a topic, set the `SubLevel` to 9.
- For multiple intercepting applications on a topic, set a lower `SubLevel` for each successive intercepting subscriber.
- You can implement a maximum of 8 intercepting applications, with subscription levels from 9 down to 2 inclusive. The final recipient of the message has a `SubLevel` of 1.

The interceptor with the highest subscription level that is equal to, or lower than, the `PubLevel` of the publication receives the publication first. Configure only one intercepting subscriber for a topic at a particular subscription level. Having multiple subscribers at a particular subscription level results in multiple copies of the publication being sent to the final set of subscribing applications.

A subscriber with a `SubLevel` of 0 is used as a catchall. It receives the publication if no final subscriber gets the message. A subscriber with `SubLevel` of 0 might be used to monitor the publications that no other subscribers received.

Programming an intercepting subscriber

Use the subscription options described in Table 27.

Table 27. Subscription options for intercepting subscribers

Subscription option	Notes
<code>MQSO_SET_CORREL_ID</code> and <code>SubCorrelId</code> set to <code>MQCI_NONE</code>	Keep the <code>CorrelId</code> of the intercepted publication the same as the original publication. Note: You cannot pass the correlation identifier of a publication in a hierarchy. The field is used by the queue manager.
<code>PubPriority</code> set to <code>MQPRI_PRIORITY_AS_PUBLISHED</code>	Keep the priority of the intercepted publication the same as the original publication.

The options in Table 27 must be used by all the intercepting subscribers. The result is that the correlation identifier and message priority are not modified from the setting of the original publisher.

When the intercepting subscriber has processed the publication, it republishes the message to the same topic at a `PubLevel` one lower than the `SubLevel` of its own subscription. If the intercepting subscriber set a `SubLevel` of 9, it republishes the message with a `PubLevel` of 8.

To republish the message correctly, several pieces of information from the original publication are required. Reuse the same **MQMD** as in the original message and set `MQPMO_PASS_ALL_CONTEXT` to ensure all information in the **MQMD** is passed on to the next subscriber. Copy the values from the message properties

shown in Table 28 into the corresponding fields of the republished message. The intercepting subscriber can change these values. Use the OR operator to add additional values to the **MQPMO.Options** field, to combine the put message options.

You must open the publication queue explicitly rather than use a managed publication queue. You cannot set **MQSO.SET_CORREL_ID** for a managed queue. You also cannot set **MQOO.SAVE_ALL_CONTEXT** on a managed queue. See the code fragments listed in “Examples.”

Table 28. MQPUT values for republished messages

Republish message using MQPUT	Information in publication message
MQOD . ObjectString	Message property MQTopicString
MQPMO . Options	Message property MQPubOptions

The final subscriber has the choice of setting its subscription options differently. For example, it might set the publication priority explicitly rather than to **MQPRI.PRIORITY_AS_PUBLISHED**. The settings of a final subscriber only affect publication from the final intercepting subscriber in the chain.

Retained publications

A retained publication must be preserved after it has been intercepted, by copying the original put-message options into the republished message.

The **MQPMO.RETAIN** option is set by the publisher. Each intercepting subscriber must transfer the **MQPubOptions** to the put-message options of the republished message as shown in Table 28. Copying the put-message options preserves the options set by the original publisher, including whether to retain the publication.

When a publication finishes its passage down the chain of intercepting subscribers, and is delivered to final subscribers, it is finally retained. New subscribers, at **SubLevel 1**, requesting the retained publication, receive it without any further interception. Subscribers at a **SubLevel** greater than 1 are not sent the retained publication. As a result, the retained publication is not modified by the chain of intercepting subscribers a second time round.

Examples

The examples are code fragments that can be combined to build an intercepting subscriber. The code is written to be brief, rather than of production quality.

The preprocessor directives in Figure 40 on page 211 define the two properties to be extracted from the publication messages that are required by the **MQINQMP MQI** call.

```

#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>
#define      MQPUBOPTIONS      (MQPTR)(char*) "MQPubOptions",\
                                0,\
                                12,\
                                MQVS_NULL_TERMINATED,\
                                MQCCSI_APPL
#define      MQTOPICSTRING     (MQPTR)(char*) "MQTopicString",\
                                0,\
                                13,\
                                MQVS_NULL_TERMINATED,\
                                MQCCSI_APPL

```

Figure 40. Preprocessor directives

Figure 41 lists the declarations used in the code fragments. Except for the highlighted terms, the declarations are standard for an IBM MQ application.

The highlighted Put and Get options are initialized to pass all context. The highlighted MQTOPICSTRING and MQPUBOPTIONS are MQCHARV initializers for property names that are defined in the preprocessor directives. The names are passed to MQINQMP.

```

int main(int argc, char **argv) {
    MQLONG Reason = MQRD_NONE;
    MQLONG CompCode = MQCC_OK;
    MQHCONN Hcon = MQHC_UNUSABLE_HCONN;
    MQCHAR QMName[49] = " ";
    MQCMHO CrtMsgHOpts = {MQCMHO_DEFAULT};
    MQHMSG Hmsg = MQHM_NONE;
    MQMD md = {MQMD_DEFAULT};
    MQHOBJ gHobj = MQHO_NONE;
    MQOD getOD = {MQOD_DEFAULT};
    MQGMO gmo = {MQGMO_DEFAULT};
    MQLONG GO_Options = MQOO_INPUT_AS_Q_DEF
        | MQOO_FAIL_IF QUIESCING
        | MQOO_SAVE_ALL_CONTEXT;
    MQLONG GC_Options = MQCO_DELETE_PURGE;
    MQHOBJ Hsub = MQHO_NONE;
    MQSD sd = {MQSD_DEFAULT};
    MQLONG SC_Options = MQCO_NONE;
    MQHOBJ pHobj = MQHO_NONE;
    MQOD putOD = {MQOD_DEFAULT};
    MQLONG PO_Options = MQOO_OUTPUT
        | MQOO_FAIL_IF QUIESCING
        | MQOO_PASS_ALL_CONTEXT;
    MQLONG PC_Options = MQCO_NONE;
    MQPMO pmo = {MQPMO_DEFAULT};
    MQIMPO InqPropOpts = {MQIMPO_DEFAULT};
    MQPD PropDesc = {MQPD_DEFAULT};
    MQLONG Type = MQTYPE_AS_SET;
    MQCHARV TopStrProp = {MQTOPICSTRING};
    MQCHARV PubOptProp = {MQPUBOPTIONS};
    MQLONG DataLength = 0;
    MQBYTE buffer[256] = "";
    MQLONG buflen = sizeof(buffer) - 1;
    MQLONG messlen = 0;
    char TopStrBuf[256] = "Initial value";
    int i = 0;
}

```

Figure 41. Declarations

Initializations that are not easily performed in declarations are shown in Figure 42 on page 212. The highlighted values require explanation.

SYSTEM.NDURABLE.MODEL.QUEUE

In this example, instead of using MQSUB to open a managed non-durable subscription, the model queue, SYSTEM.NDURABLE.MODEL.QUEUE, is used to create a temporary dynamic queue. Its handle is passed to MQSUB. By opening the queue directly you are able to save all message context and set the subscription option, MQSO_SET_CORREL_ID.

MQGMO_CURRENT_VERSION

It is important to use the current version of most of the IBM MQ structures. Fields such as gmo.MsgHandle are only available in the latest version of the control structures.

MQGMO_PROPERTIES_IN_HANDLE

The topic string and put message options set in the original publication are to be retrieved by the intercepting subscriber using message properties. An alternative would be to read the MQRFH2 structure in the message directly.

MQSO_SET_CORREL_ID

Use MQSO_SET_CORREL_ID in combination with,
memcpy(sd.SubCorrelId, MQCI_NONE, sizeof(sd.SubCorrelId));

The effect of these options is to pass on the correlation identifier. The correlation identifier set by the original publisher is placed in the correlation identifier field of the publication that is received by the intercepting subscriber. Each intercepting subscriber passes on the same correlation identifier. The final subscriber then has the option of receiving the same correlation identifier.

Note: If the publication is passed through a publish/subscribe hierarchy, the correlation identifier is never retained.

MQPRI_PRIORITY_AS_PUBLISHED

The publication is placed on the publication queue with the same message priority as it was published with.

```
strncpy(getOD.ObjectName, "SYSTEM.NDURABLE.MODEL.QUEUE",
        sizeof(getOD.ObjectName));
gmo.Version                = MQGMO_VERSION_4;
gmo.Options                = MQGMO_WAIT
                          | MQGMO_PROPERTIES_IN_HANDLE
                          | MQGMO_CONVERT;
gmo.WaitInterval          = 30000;
sd.Options                 = MQSO_CREATE
                          | MQSO_FAIL_IF QUIESCING
                          | MQSO_SET_CORREL_ID;
sd.PubPriority             = MQPRI_PRIORITY_AS_PUBLISHED;
sd.Version                 = MQSD_VERSION_1;
memcpy(sd.SubCorrelId, MQCI_NONE, sizeof(sd.SubCorrelId));
putOD.ObjectType           = MQOT_TOPIC;
putOD.ObjectString.VSPtr   = &TopStrBuf;
putOD.ObjectString.VSBufSize = sizeof(TopStrBuf);
putOD.ObjectString.VSLength = MQVS_NULL_TERMINATED;
putOD.ObjectString.VSCCSID = MQCCSI_APPL;
putOD.Version              = MQOD_VERSION_4;
pmo.Version                = MQPMO_VERSION_3;
```

Figure 42. Initializations

Figure 43 on page 213 shows the code fragment to read command-line parameters, complete the initialization, and create the intercepting subscription.

Run the program with the command,

field of `putOD.ObjectString` from an explicit length to a null terminated string is to use `printf` to output the string.

```

while (CompCode != MQCC_FAILED) {
    memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
    memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
    md.Encoding = MQENC_NATIVE;
    md.CodedCharSetId = MQCCSI_Q_MGR;
    printf("MQGET : %d seconds wait time\n", gmo.WaitInterval/1000);
    MQGET(Hcon, gHobj, &md, &gmo, buflen, buffer, &messlen,
        &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    buffer[messlen] = '\0';
    MQINQMP(Hcon, Hmsg, &InqPropOpts, &TopStrProp, &PropDesc, &Type,
        putOD.ObjectString.VSBufSize, putOD.ObjectString.VSPtr,
        &(putOD.ObjectString.VSLength), &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    memset((void *) (MQLONG) (putOD.ObjectString.VSPtr
        + putOD.ObjectString.VSLength), '\0', 1);
    putOD.ObjectString.VSLength = MQVS_NULL_TERMINATED;
    MQINQMP(Hcon, Hmsg, &InqPropOpts, &PubOptProp, &PropDesc, &Type,
        sizeof(pmo.Options), &(pmo.Options), &DataLength,
        &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    MQOPEN(Hcon, &putOD, PO_Options, &pHobj, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    printf("Republish message <%s> on topic <%s> with options %d\n",
        buffer, putOD.ObjectString.VSPtr, pmo.Options);
    MQPUT(Hcon, pHobj, &md, &pmo, messlen, buffer, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    MQCLOSE(Hcon, &pHobj, PC_Options, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
}

```

Figure 44. Intercept publication and republish

The final code fragment is shown in Figure 45.

```

} while (0);
if (CompCode == MQCC_FAILED && Reason != MQRC_NO_MSG_AVAILABLE)
    printf("MQI Call failed with reason code %d\n", Reason);
if (Hsub != MQHO_NONE)
    MQCLOSE(Hcon, &Hsub, SC_Options, &CompCode, &Reason);
if (Hcon != MQHC_UNUSABLE_HCONN)
    MQDISC(&Hcon, &CompCode, &Reason);
}

```

Figure 45. Completion

Intercepting publications and distributed publish/subscribe:

Follow a simple pattern when you deploy intercepting subscribers or Publish exits to a distributed publish/subscribe topology. Deploy intercepting subscribers on the same queue managers as publishers, and Publish exits on the same queue managers as final subscribers.

Figure 46 shows two queue managers connected in a publish subscribe cluster. A publisher creates a publication to a cluster topic at publication level 9. The numbered arrows show the sequence of steps taken by the publication as it flows to subscribers to the cluster topic. The publication is intercepted by the subscriber with SubLevel 9 and republished with PubLevel 8. It is intercepted again by a subscriber at SubLevel 8. The subscriber republishes at PubLevel 7. The proxy subscriber provided by the queue manager forwards the publication to queue manager B, where a Publish exit has been deployed in addition to a final subscriber. The publication is processed by the Publish exit before it is finally received by the final subscriber at SubLevel 1. The intercepting subscribers and the publish exit are shown with broken outlines.

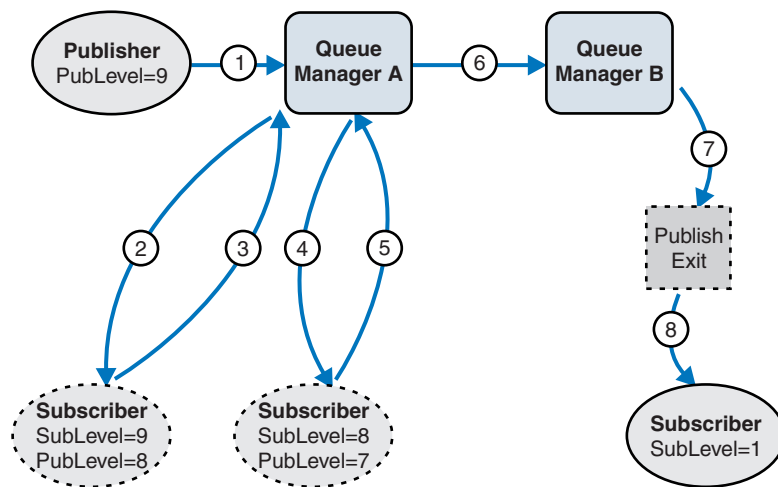


Figure 46. Interception and Publish exit in a cluster

The objective of the simple pattern is for every subscriber receiving a publication to receive the identical publication. The publication goes through the same sequence of transformations regardless of where the subscriber is connected. You probably want to avoid having the sequence of transformations varying, depending on where the publishers or final subscribers are connected. A reasonable exception would be to tailor the publication finally delivered to each individual subscriber. Use the Publish exit to do customize the publication based on the queue to which the publication is finally delivered.

You must consider carefully where to deploy intercepting subscribers and Publish exits in a distributed publish/subscribe topology. The straightforward pattern deploys intercepting subscribers to the same queue manager as the publishers, and Publish exits to the same queue managers as the final subscribers.

Anti-pattern

Figure 47 on page 216 shows how matters can go awry, if you do not follow a simple pattern. To complicate the deployment, a final subscriber is added to queue manager A and two additional intercepting subscribers are added to queue manager B.

The publication is forwarded to queue manager B at PubLevel 7, where it is intercepted by a subscriber at SubLevel 5 before being consumed by the final subscriber at SubLevel 1. The Publish exit intercepts the

publication before it is passed to both the intercepting consumer and the final consumer at queue manager B. The publication reaches the final subscriber on queue manager A without being processed by the Publish exit.

In a publish/subscribe topology, proxy subscribers subscribe at SubLevel 1, and pass on the PubLevel set by the last intercepting subscriber. In Figure 47, the result is that the publication is not intercepted by the subscriber using SubLevel 9 at queue manager B.

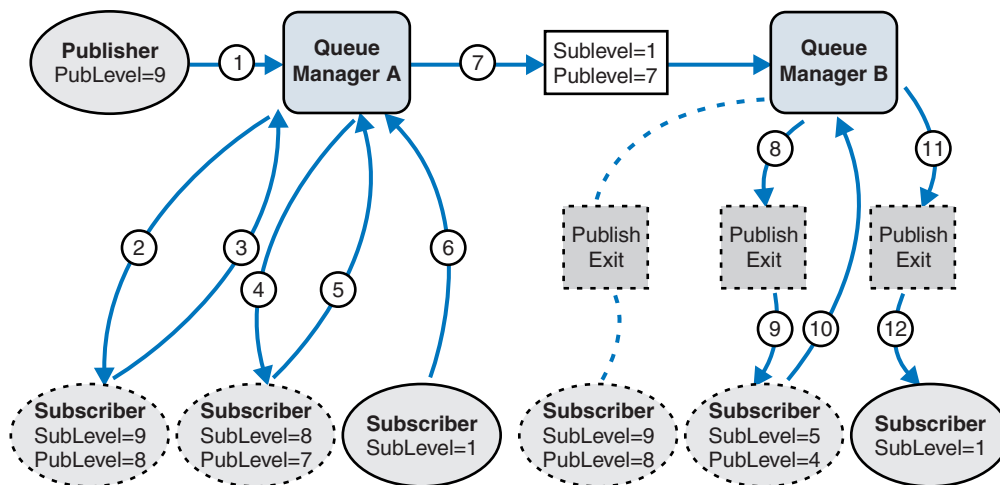


Figure 47. Complex deployment of intercepting subscribers

Publishing options:

Several options are available that control the way messages are published.

Withholding reply-to information from subscribers

If you do not want subscribers to be able to reply to publications they receive, it is possible to withhold information in the ReplyToQ and ReplyToQmgr fields of the MQMD by using the MQPMO_SUPPRESS_REPLYTO put-message option. If this option is used, the queue manager removes that information from the MQMD when it receives the publication before forwarding it to any subscribers.

This option cannot be used in combination with a report option that needs a ReplyToQ, if this is attempted the call will fail with MQRC_MISSING_REPLY_TO_Q.

Publication level

Using publication levels is a way of controlling which subscribers receive the publication. The publication level denotes the level of subscription targeted by the publication. Only subscriptions with the highest subscription level less than or equal to the publication's publication level, will receive the publication. This value must be in the range zero to nine; zero is the lowest publication level. The initial value of this field is 9. One of the uses of publication and subscription levels is to intercept publications.

Checking if a publication is not delivered to any subscribers

To check if a publication has not been delivered to any subscribers, use the MQPMO_WARN_IF_NO_SUBS_MATCHED put-message option with the MQPUT call. If a completion code of MQCC_WARNING and a reason code MQRC_NO_SUBS_MATCHED are returned by the put operation, the publication was not delivered to any subscriptions. If the MQPMO_RETAIN option is

specified on the put operation, the message is retained and delivered to any subsequently defined matching subscription. In a distributed publish/subscribe system, the MQRC_NO_SUBS_MATCHED reason code is returned only if there are no proxy subscriptions registered for the topic on the queue manager.

Subscription options:

Several options are available that control the way message subscriptions are handled.

Message persistence

Queue managers maintain the persistence of the publications they forward to subscribers as set by the publisher. The publisher sets the persistence to be one of the following options:

- 0 Nonpersistent
- 1 Persistent
- 2 Persistence as queue/topic definition

For publish/subscribe, the publisher resolves the topic object and **topicString** to a resolved topic object. If the publisher specifies Persistence as queue/topic definition, then the default persistence from the resolved topic object is set for the publication.

Retained publications

To control when retained publications are received, subscribers can use two subscription options:

Publish on request only, MQSO_PUBLICATIONS_ON_REQUEST

If you want a subscriber to have control of when it receives publications you can use the MQSO_PUBLICATIONS_ON_REQUEST subscription option. A subscriber can then control when it receives publications by using the MQSUBRQ call (specifying the Hsub handle that was returned from the original MQSUB call) to request that it is sent a topic's retained publication. Subscribers using the MQSO_PUBLICATIONS_ON_REQUEST subscription option, do not receive any non-retained publications.

If you specify MQSO_PUBLICATIONS_ON_REQUEST you must use MQSUBRQ to retrieve any publication. If you do not use MQSO_PUBLICATIONS_ON_REQUEST you get messages as they are published.

If a subscriber uses the MQSUBRQ call and uses wildcards in the subscription's topic, the subscription might match multiple topics or nodes on a topic tree, all of which with retained messages (if any exist) will be sent to the subscriber.

This option can be particularly helpful when used with durable subscriptions because a queue manager will continue to send publications to a subscriber if it subscribed durably even if that subscriber application is not running. This could lead to a buildup of messages on the subscriber queue. This build up can be avoided if the subscriber registers using the MQSO_PUBLICATIONS_ON_REQUEST option. Alternatively, you can use non-durable subscriptions if appropriate to your application to avoid a build up of unwanted messages.

If a subscription is durable and a publisher uses retained publications the subscriber application can use the MQSUBRQ call to refresh its state information after a restart. The subscriber must then refresh its state periodically using the MQSUBRQ call.

No publications will be sent as a result of the MQSUB call using this option. A durable subscription that has been resumed following disconnection will use the MQSO_PUBLICATIONS_ON_REQUEST option if the original subscription was configured to use this option.

New publications only, MQSO_NEW_PUBLICATIONS_ONLY

If a retained publication exists on a topic, any subscribers that make a subscription after the publication was made will receive a copy of that publication. If a subscriber does not want to receive any publications that were made earlier than the subscription being made, the subscriber can use the MQSO_NEW_PUBLICATIONS_ONLY subscription option.

Grouping subscriptions

Consider grouping subscriptions if you have set up a queue to receive publications and have a number of overlapping subscriptions feeding publications to the same queue. This situation is similar to the example in Overlapping subscriptions.

You can avoid receiving duplicate publications by setting the option MQSO_GROUP_SUB when you subscribe to a topic. The result is that when more than one subscription in the group matches the topic of a publication, only one subscription is responsible for placing the publication on the queue. The other subscriptions that matched the publication topic are ignored.

The subscription responsible for placing the publication on the queue is chosen on the basis that it has the longest matching topic string, before encountering any wildcards. It can be thought of as the closest matching subscription. Its properties are propagated to the publication, including whether it has the MQSO_NOT_OWN_PUBS property. If it does, no publication is delivered to the queue, even though other matching subscriptions might not have the MQSO_NOT_OWN_PUBS property.

You cannot place all your subscriptions in a single group to eliminate duplicate publications. Grouped subscriptions must fulfill these conditions:

1. None of the subscriptions are managed.
2. A group of subscriptions deliver publications to the same queue.
3. Each subscription must be at the same subscription level.
4. The publication message for each subscription in the group has the same correlation identifier.
To ensure each subscription results in a publication message with the same correlation identifier, set MQSO_SET_CORREL_ID to create your own correlation identifier in the publication, and set the same value in the **SubCorrelId** field in each subscription. Do not set **SubCorrelId** to the value MQCI_NONE.




Inquiring about and setting object attributes

Attributes are the properties that define the characteristics of an IBM MQ object.

They affect the way that a queue manager processes an object. The attributes of each type of IBM MQ object are described in detail in Attributes of objects.

Some attributes are set when the object is defined, and can be changed only by using the IBM MQ commands; an example of such an attribute is the default priority for messages put on a queue. Other attributes are affected by the operation of the queue manager and can change over time; an example is the current depth of a queue.

You can inquire about the current values of most attributes using the MQINQ call. The MQI also provides an MQSET call with which you can change some queue attributes. You cannot use the MQI calls to change the attributes of any other type of object. Instead you must use one of the following resources:

-  The MQSC facility, which is described in MQSC reference.
-  The CHGMQMx CL commands, which are described in CL commands reference for IBM i, or the MQSC facility.
-  The ALTER operator commands, or the DEFINE commands with the REPLACE option, which are described in MQSC commands.

Note: The names of the attributes of objects are shown in this documentation in the form that you use them with the MQINQ and MQSET calls. When you use IBM MQ commands to define, alter, or display the attributes, you must identify the attributes using the keywords shown in the descriptions of the commands in the topic links.

Both the MQINQ and the MQSET calls use arrays of selectors to identify those attributes that you want to inquire about or set. There is a selector for each attribute that you can work with. The selector name has a prefix, determined by the nature of the attribute:

Table 29. Prefixes for selector names

Prefix	Description
MQCA_	These selectors refer to attributes that contain character data (for example, the name of a queue).
MQIA_	These selectors refer to attributes that contain either numeric values (such as <i>CurrentQueueDepth</i> , the number of messages on a queue) or a constant value (such as <i>SyncPoint</i> , whether the queue manager supports syncpoints).

Before you use the MQINQ or MQSET calls your application must be connected to the queue manager, and you must use the MQOPEN call to open the object for setting or inquiring about attributes. These operations are described in “Connecting to and disconnecting from a queue manager” on page 91 and “Opening and closing objects” on page 100.

Use the following links to find out more about getting inquiring about and setting object attributes:

- “Inquiring about the attributes of an object” on page 220
- “Some cases where the MQINQ call fails” on page 221
- “Setting queue attributes” on page 221

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

 “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

 “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57
This information helps you to write IMS applications using IBM MQ.

Inquiring about the attributes of an object:

Use the MQINQ call to inquire about the attributes of any type of IBM MQ.

As input to this call, you must supply:

- A connection handle.
- An object handle.
- The number of selectors.
- An array of attribute selectors, each selector having the form MQCA_* or MQIA_*. Each selector represents an attribute with a value that you want to inquire about, and each selector must be valid for the type of object that the object handle represents. You can specify selectors in any order.
- The number of integer attributes that you are inquiring about. Specify zero if you are not inquiring about integer attributes.
- The length of the character attributes buffer in *CharAttrLength*. This must be at least the sum of the lengths required to hold each character attribute string. Specify zero if you are not inquiring about character attributes.

The output from MQINQ is:

- A set of integer attribute values copied into the array. The number of values is determined by *IntAttrCount*. If either *IntAttrCount* or *SelectorCount* is zero, this parameter is not used.
- The buffer in which character attributes are returned. The length of the buffer is given by the **CharAttrLength** parameter. If either *CharAttrLength* or *SelectorCount* is zero, this parameter is not used.
- A completion code. If the completion code gives a warning, this means that the call completed only partially. In this case, examine the reason code.
- A reason code. There are three partial-completion situations:
 - The selector does not apply to the queue type
 - There is not enough space allowed for integer attributes
 - There is not enough space allowed for character attributes

If more than one of these situations arise, the first one that applies is returned.

If you open a queue for output or inquire and it resolves to a non-local cluster queue you can only inquire the queue name, queue type, and common attributes. The values of the common attributes are those of the chosen queue if MQOO_BIND_ON_OPEN was used. The values are those of an arbitrary one of the possible cluster queues if either MQOO_BIND_NOT_FIXED or MQOO_BIND_ON_GROUP was used or MQOO_BIND_AS_Q_DEF was used and the **DefBind** queue attribute was MQBND_BIND_NOT_FIXED. See “MQOPEN and clusters” on page 254 and MQOPEN for more information.

Note: The values returned by the call are a snapshot of the selected attributes. The attributes can change before your program acts on the returned values.

There is a description of the MQINQ call in MQINQ.

Some cases where the MQINQ call fails:

If you open an alias to inquire about its attributes, you are returned the attributes of the alias queue (the IBM MQ object used to access another queue), not those of the base queue.

However, the definition of the base queue to which the alias resolves is also opened by the queue manager, and if another program changes the usage of the base queue in the interval between your MQOPEN and MQINQ calls, your MQINQ call fails and returns the MQRC_OBJECT_CHANGED reason code. The call also fails if the attributes of the alias queue object are changed.

Similarly, when you open a remote queue to inquire about its attributes, you are returned the attributes of the local definition of the remote queue only.

If you specify one or more selectors that are not valid for the type of queue attributes upon which you are inquiring, the MQINQ call completes with a warning and sets the output as follows:

- For integer attributes, the corresponding elements of *IntAttrs* are set to MQIAV_NOT_APPLICABLE.
- For character attributes, the corresponding portions of the *CharAttrs* string are set to asterisks.

If you specify one or more selectors that are not valid for the type of object attributes upon which you are inquiring, the MQINQ call fails and returns the MQRC_SELECTOR_ERROR reason code.

You cannot call MQINQ to look at a model queue; use either the MQSC facility or the commands available on your platform.

Setting queue attributes:

Use this information to learn how to set queue attributes using the MQSET call.

You can set only the following queue attributes using the MQSET call:

- *InhibitGet* (but not for remote queues)
- *DistList* (not on z/OS)
- *InhibitPut*
- *TriggerControl*
- *TriggerType*
- *TriggerDepth*
- *TriggerMsgPriority*
- *TriggerData*

The MQSET call has the same parameters as the MQINQ call. However, for MQSET, all parameters except the completion code and reason code are input parameters. There are no partial-completion situations.

Note: You cannot use the MQI to set the attributes of IBM MQ objects other than locally defined queues.

For more details about the MQSET call, see MQSET.

Committing and backing out units of work

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

The following terms are used in this topic:

- Commit
- Back out
- Syncpoint coordination
- Syncpoint
- Unit of work
- Single-phase commit
- Two-phase commit

If you are familiar with these transaction processing terms, you can skip to “Syncpoint considerations in IBM MQ applications” on page 223.

Commit and back out


When a program puts a message on a queue within a unit of work, that message is made visible to other programs only when the program commits the unit of work. To commit a unit of work, all updates must be successful to preserve data integrity. If the program detects an error and decides that the put operation is not permanent, it can back out the unit of work. When a program performs a backout, IBM MQ restores the queue by removing the messages that were put on the queue by that unit of work. The way in which the program performs the commit and back out operations depends on the environment in which the program is running.


Similarly, when a program gets a message from a queue within a unit of work, that message remains on the queue until the program commits the unit of work, but the message is not available to be retrieved by other programs. The message is permanently deleted from the queue when the program commits the unit of work. If the program backs out the unit of work, IBM MQ restores the queue by making the messages available to be retrieved by other programs.

Syncpoint coordination, syncpoint, unit of work

Syncpoint coordination is the process by which units of work are either committed or backed out with data integrity.

The decision to commit or back out the changes is taken, in the simplest case, at the end of a transaction. However, it can be more useful for an application to synchronize data changes at other logical points within a transaction. These logical points are called *syncpoints* (or *synchronization points*) and the period of processing a set of updates between two syncpoints is called a *unit of work*. Several MQGET calls and MQPUT calls can be part of a single unit of work.

 On z/OS, the maximum number of messages within a unit of work can be controlled by the DEFINE MAXSMGS command.

 On Multiplatforms, the maximum number of messages within a unit of work can be controlled by the MAXUMSGS attribute of the ALTER QMGR command.

Single-phase commit

A *single-phase commit* process is one in which a program can commit updates to a queue without coordinating its changes with other resource managers.




Two-phase commit

A *two-phase commit* process is one in which updates that a program has made to IBM MQ queues can be coordinated with updates to other resources (for example, databases under the control of Db2). Under such a process, updates to all resources are committed or backed out together.

To help handle units of work, IBM MQ provides the **BackoutCount** attribute. This is incremented each time that a message within a unit of work is backed out. If the message repeatedly causes

the unit of work to abnormally end, the value of the *BackoutCount* finally exceeds that of the *BackoutThreshold*. This value is set when the queue is defined. In this situation, the application can remove the message from the unit of work and put it onto another queue, as defined in *BackoutRequeueQName*. When the message is moved, the unit of work can commit.

Use the following links to find out more about committing and backing out units of work:

- “Syncpoint considerations in IBM MQ applications”
-  “Syncpoints in IBM MQ for z/OS applications” on page 225
-  “Syncpoints in CICS for IBM i applications” on page 227
- “Syncpoints in IBM MQ for IBM i, UNIX, Linux, and Windows” on page 228
-  “Interfaces to the IBM i external syncpoint manager” on page 233

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

 “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.







 “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57





This information helps you to write IMS applications using IBM MQ.

Syncpoint considerations in IBM MQ applications:





Use this information to learn about using syncpoints in IBM MQ applications.

Two-phase commit is supported by the following environments:

-  IBM MQ for AIX
-  IBM MQ for IBM i
-  IBM MQ for HP-UX
-  IBM MQ for Linux
-  IBM MQ for Solaris
-  IBM MQ for Windows

-  CICS Transaction Server for z/OS
-  TXSeries
-  IMS/ESA®
-  z/OS batch with RRS
- Other external coordinators using the X/Open XA interface



Single-phase commit is supported by the following environments:

-  IBM MQ for IBM i
-  IBM MQ on UNIX
-  IBM MQ for Windows
-  z/OS batch

For further information about external interfaces, see “Interfaces to external syncpoint managers on Multiplatforms” on page 231, and the XA documentation *CAE Specification Distributed Transaction Processing: The XA Specification*, published by The Open Group. Transaction managers (such as CICS, IMS, Encina, and Tuxedo) can participate in two-phase commit, coordinated with other recoverable resources. This means that the queuing functions provided by IBM MQ can be brought within the scope of a unit of work, managed by the transaction manager.

Samples shipped with IBM MQ show IBM MQ coordinating XA-compliant databases. For further information about these samples, see “Using the IBM MQ sample procedural programs” on page 457.

In your IBM MQ application, you can specify on every put and get call whether you want the call to be under syncpoint control. To make a put operation operate under syncpoint control, use the MQPMO_SYNCPOINT value in the *Options* field of the MQPMO structure when you call MQPUT. For a get operation, use the MQGMO_SYNCPOINT value in the *Options* field of the MQGMO structure. If you do not explicitly choose an option, the default action depends on the platform:

-  The syncpoint control default is NO.
-  The syncpoint control default is YES.

When an MQPUT1 call is issued with MQPMO_SYNCPOINT, the default behavior changes, so that the put operation is completed asynchronously. This might cause a change in the behavior of some applications that rely on certain fields in the MQOD and MQMD structures being returned, but which now contain undefined values. An application can specify MQPMO_SYNC_RESPONSE to ensure that the put operation is performed synchronously and that all of the appropriate field values are completed.

When your application receives an MQRC_BACKED_OUT reason code in response to an MQPUT or MQGET under syncpoint, the application should normally back out the current transaction using MQBACK and then, if appropriate, try the entire transaction again. If the application receives MQRC_BACKED_OUT in response to an MQCMIT or MQDISC call, it does not need to call MQBACK.

Every time an MQGET call is backed out, the *BackoutCount* field of the MQMD structure of the affected message is incremented. A high *BackoutCount* indicates a message that has been repeatedly backed out. This might indicate a problem with this message, which you should investigate. See *BackoutCount* for details of *BackoutCount*.

Except on z/OS batch with RRS, if a program issues the MQDISC call while there are uncommitted requests, an implicit syncpoint occurs. If the program ends abnormally, an implicit backout occurs.

▶ **z/OS** On z/OS, an implicit syncpoint also occurs if the program ends normally without first calling MQDISC. The program is deemed to have ended normally if the TCB connected to MQ ends normally. When running under z/OS UNIX System Services and Language Environment (LE), default condition handling is invoked for abends or signals. The LE condition handlers process the error condition and the TCB ends normally. Under these conditions MQ commits the unit of work. For more information, see Introduction to Language Environment Condition Handling and Customizing Language Environment run-time options.

▶ **z/OS** For IBM MQ for z/OS programs, you can use the MQGMO_MARK_SKIP_BACKOUT option to specify that a message must not be backed out if backout occurs (in order to avoid an *MQGET-error-backout* loop). For information about using this option, see “Skipping backout” on page 161.

Changes to queue attributes (either by the MQSET call or by commands) are not affected by the committing or backing out of units of work.

Syncpoints in IBM MQ for z/OS applications: ▶ **z/OS**

This topic explains how to use syncpoints in transaction manager (CICS and IMS) and batch applications.

Syncpoints in CICS Transaction Server for z/OS applications: ▶ **z/OS**

In a CICS application you establish a syncpoint by using the EXEC CICS SYNCPOINT command.

To back out all changes to the previous syncpoint, you can use the EXEC CICS SYNCPOINT ROLLBACK command. For more information, see the *CICS Application Programming Reference*.

If other recoverable resources are involved in the unit of work, the queue manager (in conjunction with the CICS syncpoint manager) participates in a two-phase commit protocol; otherwise, the queue manager performs a single-phase commit process.

If a CICS application issues the MQDISC call, no implicit syncpoint is taken. If the application closes down normally, any open queues are closed and an implicit commit occurs. If the application closes down abnormally, any open queues are closed and an implicit backout occurs.

Syncpoints in IMS applications: ▶ **z/OS**

In an IMS application, establish a syncpoint by using IMS calls such as GU (get unique) to the IOPCB and CHKP (checkpoint).

To back out all changes since the previous checkpoint, you can use the IMS ROLB (rollback) call. For more information, see the IMS documentation.

The queue manager (in conjunction with the IMS syncpoint manager) participates in a two-phase commit protocol if other recoverable resources are also involved in the unit of work.

All open handles are closed by the IMS adapter at a syncpoint (except in a batch or non-message driven BMP environment). This is because a different user could initiate the next unit of work and IBM MQ security checking is performed when the MQCONN, MQCONNX, and MQOPEN calls are made, not when the MQPUT or MQGET calls are made.

However, in a Wait-for-Input (WFI) or pseudo Wait-for-Input (PWFI) environment IMS does not notify IBM MQ to close the handles until either the next message arrives or a QC status code is returned to the application. If the application is waiting in the IMS region and any of these handles belong to triggered

queues, triggering will not occur because the queues are open. For this reason applications running in a WFI or PWFI environment should explicitly MQCLOSE the queue handles before doing the GU to the IOPCB for the next message.

If an IMS application (either a BMP or an MPP) issues the MQDISC call, open queues are closed but no implicit syncpoint is taken. If the application closes down normally, any open queues are closed and an implicit commit occurs. If the application closes down abnormally, any open queues are closed and an implicit backout occurs.

Syncpoints in z/OS batch applications: 

For batch applications, you can use the IBM MQ syncpoint management calls: MQCMIT and MQBACK. For compatibility with earlier versions, CSQBCMT and CSQBBAK are available as synonyms.

Note: If you need to commit or back out updates to resources managed by different resource managers, such as IBM MQ and Db2, within a single unit of work you can use RRS. For further information see “Transaction management and recoverable resource manager services.”

Committing changes using the MQCMIT call

As input, you must supply the connection handle (*Hconn*) that is returned by the MQCONN or MQCONNX call.

The output from MQCMIT is a completion code and a reason code. The call completes with a warning if the syncpoint was completed but the queue manager backed out the put and get operations since the previous syncpoint.

Successful completion of the MQCMIT call indicates to the queue manager that the application has reached a syncpoint and that all put and get operations made since the previous syncpoint have been made permanent.

Not all failure responses mean that the MQCMIT did not complete. For example, the application can receive MQRC_CONNECTION_BROKEN.

There is a description of the MQCMIT call in MQCMIT.

Backing out changes using the MQBACK call

As input, you must supply a connection handle (*Hconn*). Use the handle that is returned by the MQCONN or MQCONNX call.

The output from MQBACK is a completion code and a reason code.

The output indicates to the queue manager that the application has reached a syncpoint and that all gets and puts that have been made since the last syncpoint have been backed out.

There is a description of the MQBACK call in MQBACK.

Transaction management and recoverable resource manager services

Transaction management and recoverable resource manager services (RRS) is a z/OS facility to provide two-phase syncpoint support across participating resource managers.

An application can update recoverable resources managed by various z/OS resource managers such as IBM MQ and Db2, and then commit or back out these updates as a single unit of work. RRS provides the

necessary unit-of-work status logging during normal execution, coordinates the syncpoint processing, and provides appropriate unit-of-work status information during subsystem restart.

IBM MQ for z/OS RRS participant support enables IBM MQ applications in the batch, TSO, and Db2 stored procedure environments to update both IBM MQ and non-IBM MQ resources (for example, Db2) within a single logical unit of work. For information about RRS participant support, see *MVS Programming: Resource Recovery*.

Your IBM MQ application can use either MQCMIT and MQBACK or the equivalent RRS calls, SRRCMIT and SRRBACK. See “The RRS batch adapter” on page 261 for more information.

RRS availability

If RRS is not active on your z/OS system, any IBM MQ call issued from a program linked with either RRS stub (CSQBRSTB or CSQBRRSI) returns MQRC_ENVIRONMENT_ERROR.

Db2 stored procedures

If you use Db2 stored procedures with RRS, be aware of the following:

- Db2 stored procedures that use RRS must be managed by workload manager (WLM-managed).
- If a Db2-managed stored procedure contains IBM MQ calls, and it is linked with either RRS stub (CSQBRSTB or CSQBRRSI), the MQCONN or MQCONNX call returns MQRC_ENVIRONMENT_ERROR.
- If a WLM-managed stored procedure contains IBM MQ calls, and is linked with a non-RRS stub, the MQCONN or MQCONNX call returns MQRC_ENVIRONMENT_ERROR, unless it is the first IBM MQ call executed since the stored procedure address space started.
- If your Db2 stored procedure contains IBM MQ calls and is linked with a non-RRS stub, IBM MQ resources updated in that stored procedure are not committed until the stored procedure address space ends, or until a subsequent stored procedure does an MQCMIT (using an IBM MQ Batch/TSO stub).
- Multiple copies of the same stored procedure can execute concurrently in the same address space. Ensure that your program is coded in a reentrant manner if you want Db2 to use a single copy of your stored procedure. Otherwise you might receive MQRC_HCONN_ERROR on any IBM MQ call in your program.
- Do not code MQCMIT or MQBACK in a WLM-managed Db2 stored procedure.
- Design all programs to run in Language Environment (LE).

Syncpoints in CICS for IBM i applications:

IBM MQ for IBM i participates in CICS for IBM i units of work. You can use the MQI within a CICS for IBM i application to put and get messages inside the current unit of work.

You can use the EXEC CICS SYNCPOINT command to establish a syncpoint that includes the IBM MQ for IBM i operations. To back out all changes up to the previous syncpoint, you can use the EXEC CICS SYNCPOINT ROLLBACK command.


If you use MQPUT, MQPUT1, or MQGET with the MQPMO_SYNCPOINT, or MQGMO_SYNCPOINT, option set in a CICS for IBM i application, you cannot log off CICS for IBM i until IBM MQ for IBM i has removed its registration as an API commitment resource. Commit or back out any pending put or get operations before you disconnect from the queue manager. This allows you to log off CICS for IBM i.

Syncpoints in IBM MQ for IBM i, UNIX, Linux, and Windows:

Syncpoint support operates on two types of units of work: local and global.

A *local* unit of work is one in which the only resources updated are those of the IBM MQ queue manager. Here syncpoint coordination is provided by the queue manager itself using a single-phase commit procedure.


A *global* unit of work is one in which resources belonging to other resource managers, such as databases, are also updated. IBM MQ can coordinate such units of work itself. They can also be coordinated by an external commitment controller. For example:

- Another transaction manager
-  The IBM i commitment controller

For full integrity, use a two-phase commit procedure. Two-phase commit can be provided by XA-compliant transaction managers and databases. For example:

- TXSeries
- UDB
-  the IBM i commitment controller

 IBM MQ products can coordinate global units of work using a two-phase commit process.

 IBM MQ for IBM i can act as a resource manager for global units of work within a WebSphere Application Server environment, but cannot act as a transaction manager.



Implicit syncpoint

When putting persistent messages, IBM MQ is optimized for putting persistent messages under syncpoint. Multiple applications putting persistent messages to the same queue perform better if those applications use syncpoint. This is because there is less contention for the queue, if syncpoint is used to put persistent messages.

ImplSyncOpenOutput adds an implicit syncpoint when applications put persistent messages outside of syncpoint. This provides a performance improvement, without applications being aware of the implicit syncpoint.

Implicit syncpoint only provides a performance boost when there are multiple applications putting to the queue, because it reduces contention for the queue. So, **ImplSyncOpenOutput** specifies the minimum number of applications that have a queue open for output before an implicit syncpoint is added. The default value is 2. This means, that if you do not specify **ImplSyncOpenOutput**, implicit syncpoint is only added if multiple applications are putting to the queue.

See Tuning parameters for more information.

Local units of work:

Units of work that involve only the queue manager are called *local* units of work. Syncpoint coordination is provided by the queue manager itself (internal coordination) using a single-phase commit process.

To start a local unit of work, the application issues MQGET, MQPUT, or MQPUT1 requests specifying the appropriate syncpoint option. The unit of work is committed using MQCMIT or rolled back using MQBACK. However, the unit of work also ends when the connection between the application and the queue manager is broken, intentionally or unintentionally.

If an application disconnects (MQDISC) from a queue manager while a global unit of work coordinated by IBM MQ is still active, an attempt is made to commit the unit of work. If, however, the application terminates without disconnecting, the unit of work is rolled back as the application is deemed to have terminated abnormally.

Global units of work:

Use global units of work when you also need to include updates to resources belonging to other resource managers.

Here the coordination can be internal or external to the queue manager:

Internal syncpoint coordination

Queue manager coordination of global units of work is not supported by IBM MQ for IBM i or IBM MQ for z/OS. It is not supported in an IBM MQ MQI client environment.

Here, IBM MQ does the coordination. To start a global unit of work, the application issues the MQBEGIN call.

As input to the MQBEGIN call, you must supply the connection handle (*Hconn*) that is returned by the MQCONN or MQCONNX call. This handle represents the connection to the IBM MQ queue manager.

The application issues MQGET, MQPUT, or MQPUT1 requests specifying the appropriate syncpoint option. This means that you can use MQBEGIN to initiate a global unit of work that updates local resources, resources belonging to other resource managers, or both. Updates made to resources belonging to other resource managers are made using the API of that resource manager. However, you cannot use the MQI to update queues that belong to other queue managers. Issue MQCMIT or MQBACK before starting further units of work (local or global).

The global unit of work is committed using MQCMIT; this initiates a two-phase commit of all the resource managers involved in the unit of work. A two-phase commit process is used whereby resource managers (for example, XA-compliant database managers such as Db2, Oracle, and Sybase) are first all asked to prepare to commit. Only if all are prepared are they asked to commit. If any resource manager signals that it cannot commit, each is asked to back out instead. Alternatively, you can use MQBACK to roll back the updates of all the resource managers.

If an application disconnects (MQDISC) while a global unit of work is still active, the unit of work is committed. If, however, the application terminates without disconnecting, the unit of work is rolled back as the application is deemed to have terminated abnormally.

The output from MQBEGIN is a completion code and a reason code.

When you use MQBEGIN to start a global unit of work, all the external resource managers that have been configured with the queue manager are included. However, the call starts a unit of work but completes with a warning if:

- There are no participating resource managers (that is, no resource managers have been configured with the queue manager)

or

- One or more resource managers are not available.

In these cases, the unit of work must include updates to only those resource managers that were available when the unit of work was started.

If one of the resource managers cannot commit its updates, all the resource managers are instructed to roll back their updates, and MQCMIT completes with a warning. In unusual circumstances (typically, operator intervention), an MQCMIT call might fail if some resource managers commit their updates but others roll them back; the work is deemed to have completed with a *mixed* outcome. Such occurrences are diagnosed in the error log of the queue manager so that remedial action may be taken.

An MQCMIT of a global unit of work succeeds if all the resource managers involved commit their updates.

For a description of the MQBEGIN call, see MQBEGIN.

External syncpoint coordination

This occurs when a syncpoint coordinator other than IBM MQ has been selected; for example, CICS, Encina, or Tuxedo.

In this situation, IBM MQ on UNIX and Linux systems and IBM MQ for Windows register their interest in the outcome of the unit of work with the syncpoint coordinator so that they can commit or roll back any uncommitted get or put operations as required. The external syncpoint coordinator determines whether one- or two-phase commitment protocols are provided.

When you use an external coordinator, MQCMIT, MQBACK, and MQBEGIN cannot be issued. Calls to these functions fail with the reason code MQRC_ENVIRONMENT_ERROR.


The way in which an externally coordinated unit of work is started depends on the programming interface provided by the syncpoint coordinator. An explicit call might be required. If an explicit call is required, and you issue an MQPUT call specifying the MQPMO_SYNCPOINT option when a unit of work is not started, the completion code MQRC_SYNCPOINT_NOT_AVAILABLE is returned.

The scope of the unit of work is determined by the syncpoint coordinator. The state of the connection between the application and the queue manager affects the success or failure of MQI calls that an application issues, not the state of the unit of work. An application can, for example, disconnect and reconnect to a queue manager during an active unit of work and perform further MQGET and MQPUT operations inside the same unit of work. This is known as a pending disconnect.

You can use IBM MQ API calls in CICS programs, whether you choose to use the XA abilities of CICS. If you do not use XA, then the puts and gets of messages to and from queues will not be managed within CICS atomic units of work. One reason for choosing this method is that the overall consistency of the unit of work is not important to you.

If the integrity of your units of work is important to you, then you must use XA. When you use XA, CICS uses a two-phase commit protocol to ensure all resources within the unit of work are updated together.

For more information about setting up transactional support, see Transactional support scenarios, and also TXSeries CICS documentation, for example, *TXSeries for Multiplatforms CICS Administration Guide for Open Systems*.

Implicit syncpoint: 

IBM MQ Version 9.0.5 adds an implicit syncpoint for persistent messages put outside of syncpoint.

When putting persistent messages, IBM MQ is optimized for putting persistent messages under syncpoint. Multiple applications concurrently putting persistent messages to the same queue typically perform better if those applications use syncpoint. This is because the locking strategy of IBM MQ is more efficient if syncpoint is used when putting persistent messages.

The **ImplSyncOpenOutput** parameter in the `qm.ini` file, controls whether an implicit syncpoint can be added when applications put persistent messages outside of syncpoint. This can provide a performance improvement, without applications being aware of the implicit syncpoint.

Implicit syncpoint only provides a performance boost when there are multiple applications concurrently putting to the queue, because it reduces lock contention. **ImplSyncOpenOutput** specifies the minimum number of applications that have a queue open for output before an implicit syncpoint can be added. The default value is 2. This means, that if you do not explicitly specify **ImplSyncOpenOutput**, implicit syncpoint is only added if multiple applications are putting to the queue.


If you add an implicit syncpoint, statistics reflect that happening, and you might see a transaction output from **runmqsc display conn**.

Set **ImplSyncOpenOutput=OFF** if you never want an implicit syncpoint added.

See Tuning parameters for more information.

Interfaces to external syncpoint managers on Multiplatforms:

IBM MQ for Multiplatforms supports coordination of transactions by external syncpoint managers that use the X/Open XA interface.

Some XA transaction managers (TXSeries) require that each XA resource manager supplies its name. This is the string called name in the XA switch structure. The resource manager for IBM MQ on UNIX, Linux, and Windows is named `MQSeries_XA_RMI`.  For IBM i, the resource manager name is `MQSeries XA RMI`. For further details on XA interfaces refer to the XA documentation *CAE Specification Distributed Transaction Processing: The XA Specification*, published by The Open Group.

In an XA configuration, IBM MQ for Multiplatforms fulfills the role of an XA resource manager. An XA syncpoint coordinator can manage a set of XA resource managers, and synchronize the commit or backout of transactions in both resource managers. This is how it works for a statically-registered resource manager:

1. An application notifies the syncpoint coordinator that it wants to start a transaction.
2. The syncpoint coordinator issues a call to any resource managers that it knows of, to notify them of the current transaction.
3. The application issues calls to update the resources managed by the resource managers associated with the current transaction.
4. The application requests that the syncpoint coordinator either commit or roll back the transaction.
5. The syncpoint coordinator issues calls to each resource manager using two-phase commit protocols to complete the transaction as requested.

The XA specification requires each resource manager to provide a structure called an XA Switch. This structure declares the capabilities of the resource manager, and the functions that are to be called by the syncpoint coordinator.

There are two versions of this structure:

Table 30. Versions of XA Switch

Version	Description
MQRMIXASwitch	Static XA resource management
MQRMIXASwitchDynamic	Dynamic XA resource management

For a list of the libraries containing this structure see The IBM MQ XA switch structure.


The method that must be used to link them to an XA syncpoint coordinator is defined by the coordinator; consult the documentation provided by that coordinator to determine how to enable IBM MQ to cooperate with your XA syncpoint coordinator.

The *xa_info* structure that is passed on any *xa_open* call by the syncpoint coordinator can be the name of the queue manager that is to be administered. This takes the same form as the queue manager name passed to MQCONN or MQCONNX, and can be blank if the default queue manager is to be used. However, you can use the two extra parameters TPM and AXLIB

TPM allows you to specify to IBM MQ the transaction manager name, for example, CICS. AXLIB allows you to specify the actual library name in the transaction manager where the XA AX entry points are located.

If you use either of these parameters or a non default queue manager you must specify the queue manager name using the QMNAME parameter. For further information see The CHANNEL, TRPTYPE, CONNAME, and QMNAME parameters of the *xa_open* string.

Restrictions

1. Global units of work are not allowed with a shared Hconn (as described in “Shared (thread independent) connections with MQCONNX” on page 97.
2.  IBM MQ for IBM i does not support dynamic registration of XA resource managers. The only transaction manager supported is WebSphere Application Server.
3. On Windows systems, all functions declared in the XA switch are declared as `_cdecl` functions.
4. An external syncpoint coordinator can administer only one queue manager at a time. This is because the coordinator has an effective connection to each queue manager, and is therefore subject to the rule that only one connection is allowed at a time.

Note: Note: A JMS client application (CLIENT JEE application) running in a JEE server does not have this restriction, so a single JEE server-managed transaction can coordinate multiple queue managers in the same transaction. However, a JMS server application, running in bindings mode, is still subject to the rule that only one connection is allowed at a time.

5. All applications that are run using the syncpoint coordinator can connect only to the queue manager that is administered by the coordinator because they are already effectively connected to that queue manager. They must issue MQCONN or MQCONNX to obtain a connection handle and must issue MQDISC before they exit. Alternatively, they can use the exit UE014015 for TXSeries CICS.

Interfaces to the IBM i external syncpoint manager:

IBM i

IBM MQ for IBM i can use native IBM i commitment control as an external syncpoint coordinator.

Thread-independent (shared) connections are not allowed with commitment control. See the *IBM i Programming: Backup and Recovery Guide, SC21-8079* for more information about the commitment control capabilities of IBM i.

To start the IBM i commitment control facilities, use the STRCMTCTL system command. To end commitment control, use the ENDCMTCTL system command.

Note: The default value of *Commitment definition scope* is *ACTGRP. This must be defined as *JOB for IBM MQ for IBM i. For example:

```
STRCMTCTL LCKLVL(*ALL) CMTSCOPE(*JOB)
```

IBM MQ for IBM i can also perform local units of work containing only updates to IBM MQ resources. The choice between local units of work and participation in global units of work coordinated by IBM i is made in each application when the application calls MQPUT, MQPUT1, or MQGET, specifying MQPMO_SYNCPOINT or MQGMO_SYNCPOINT, or MQBEGIN. If commitment control is not active when the first such call is issued, IBM MQ starts a local unit of work and all further units of work for this connection to IBM MQ also use local units of work, regardless of whether commitment control is then started. To commit a local unit of work, use MQCMIT. To back out a local unit of work, use MQBACK. The IBM i commit and rollback calls such as the CL command COMMIT have no effect on IBM MQ local units of work.

If you want to use IBM MQ for IBM i with native IBM i commitment control as an external syncpoint coordinator, ensure that any job with commitment control is active and that you are using IBM MQ in a single-threaded job. If you call MQPUT, MQPUT1, or MQGET, specifying MQPMO_SYNCPOINT or MQGMO_SYNCPOINT, in a multithreaded job in which commitment control has been started, the call fails with a reason code of MQRC_SYNCPOINT_NOT_AVAILABLE.

It is possible to use local units of work and the MQCMIT and MQBACK calls in a multithreaded job.

If you call MQPUT, MQPUT1, or MQGET, specifying MQPMO_SYNCPOINT or MQGMO_SYNCPOINT, after starting commitment control, IBM MQ for IBM i adds itself as an API commitment resource to the commitment definition. This is typically the first such call in a job. While there are any API commitment resources registered under a particular commitment definition, you cannot end commitment control for that definition.

IBM MQ for IBM i removes its registration as an API commitment resource when you disconnect from the queue manager, if there are no pending MQI operations in the current unit of work.

If you disconnect from the queue manager while there are pending MQPUT, MQPUT1, or MQGET operations in the current unit of work, IBM MQ for IBM i remains registered as an API commitment resource so that it is notified of the next commit or rollback. When the next syncpoint is reached, IBM MQ for IBM i commits or rolls back the changes as required. An application can disconnect and reconnect to a queue manager during an active unit of work and perform further MQGET and MQPUT operations inside the same unit of work (this is a pending disconnect).

If you attempt to issue an ENDCMTCTL system command for that commitment definition, message CPF8355 is issued, indicating that pending changes were active. This message also appears in the job log when the job ends. To avoid this, commit or roll back all pending IBM MQ for IBM i operations, and disconnect from the queue manager. Thus, using COMMIT or ROLLBACK commands before ENDCMTCTL enables end-commitment control to complete successfully.

When you use IBM i commitment control as an external syncpoint coordinator, you cannot issue MQCMIT, MQBACK, and MQBEGIN calls. Calls to these functions fail with the reason code MQRC_ENVIRONMENT_ERROR.

To commit or roll back (that is, to back out) your unit of work, use one of the programming languages that supports the commitment control. For example:

- CL commands: COMMIT and ROLLBACK
- ILE C Programming Functions: `_Rcommit` and `_Rrollback`
- ILE RPG: COMMIT and ROLBK
- COBOL/400®: COMMIT and ROLLBACK

When you use IBM i commitment control as an external syncpoint coordinator with IBM MQ for IBM i, IBM i performs a two-phase commit protocol in which IBM MQ participates. Because each unit of work is committed in two phases, the queue manager might become unavailable for the second phase after having voted to commit in the first phase. This can happen, for example, if the queue manager's internal jobs are ended. In this situation, the job log performing the commit contains message CPF835F indicating that a commit or rollback operation failed. The messages preceding this indicate the cause of the problem, whether it occurred during a commit or rollback operation, and also the logical unit of work ID (LUWID) for the failed unit of work.

If the problem was caused by the failure of the IBM MQ API commitment resource during the commit or rollback of a prepared unit of work, you can use the WRKMQMTRN command to complete the operation and restore the integrity of the transaction. The command requires that you know the LUWID of the unit of work to commit and back out.

Starting IBM MQ applications using triggers

Learn about triggers and how to start IBM MQ applications using triggers.

Some IBM MQ applications that serve queues run continuously, so they are always available to retrieve messages that arrive on the queues. However, you might not want this when the number of messages arriving on the queues is unpredictable. In this case, applications could be consuming system resources even when there are no messages to retrieve.

IBM MQ provides a facility that enables an application to be started automatically when there are messages available to retrieve. This facility is known as *triggering*.

For information about triggering channels see [Triggering channels](#).

What is triggering?

The queue manager defines certain conditions as constituting *trigger events*.

If triggering is enabled for a queue and a trigger event occurs, the queue manager sends a *trigger message* to a queue called an *initiation queue*. The presence of the trigger message on the initiation queue indicates that a trigger event has occurred.

Trigger messages generated by the queue manager are not persistent. This reduces logging (resulting in improving performance), and minimizing duplicates during restart, so improving restart time.

The program that processes the initiation queue is called a *trigger-monitor application*, and its function is to read the trigger message and take appropriate action, based on the information contained in the trigger message. Typically this action is to start some other application to process the queue that generated the trigger message. From the point of view of the queue manager, there is nothing special about the trigger-monitor application; it is simply another application that reads messages from a queue (the initiation queue).

If triggering is enabled for a queue, you can create a *process-definition object* associated with it. This object contains information about the application that processes the message that caused the trigger event. If the process definition object is created, the queue manager extracts this information and places it in the trigger message, for use by the trigger-monitor application. The name of the process definition associated with a queue is given by the *ProcessName* local-queue attribute. Each queue can specify a different process definition, or several queues can share the same process definition.

If you want to trigger the start of a channel, you do not need to define a process definition object. The transmission queue definition is used instead.

Triggering is supported by IBM MQ clients running on UNIX, Linux, and Windows. An application running in a client environment is the same as one running in a full IBM MQ environment, except that you link it with the client libraries. However the trigger monitor and the application to be started must both be in the same environment.

Triggering involves:

Application queue

An *application queue* is a local queue that, when it has triggering set on and when the conditions are met, requires that trigger messages are written.

Process definition

An application queue can have a *process definition object* associated with it that holds details of the application that will get messages from the application queue. (See Attributes for process definitions for a list of attributes.)

Remember that if you want a trigger to start a channel, you do not need to define a process definition object.

Transmission queue

You need a transmission queue if you want a trigger to start a channel.

For a transmission queue on any platform other than Linux, the *TriggerData* attribute of the transmission queue can specify the name of the channel to be started. This can replace the process definition for triggering channels, but is used only when a process definition is not created.

Trigger event

A *trigger event* is an event that causes a trigger message to be generated by the queue manager. This is typically a message arriving on an application queue, but it can also occur at other times. For example, see “Conditions for a trigger event” on page 241.

IBM MQ has a range of options to allow you to control the conditions that cause a trigger event (see “Controlling trigger events” on page 245).

Trigger message


The queue manager creates a *trigger message* when it recognizes a trigger event. It copies into the trigger message information about the application to be started. This information comes from the application queue and the process definition object associated with the application queue.

Trigger messages have a fixed format (see “Format of trigger messages” on page 252).

Initiation queue

An *initiation queue* is a local queue on which the queue manager puts trigger messages. Note that an initiation queue cannot be an alias queue or a model queue.

A queue manager can own more than one initiation queue, and each one is associated with one or more application queues.

 A shared queue, a local queue accessible by queue managers in a queue-sharing group, can be an initiation queue on IBM MQ for z/OS.

Trigger monitor

A *trigger monitor* is a continuously running program that serves one or more initiation queues. When a trigger message arrives on an initiation queue, the trigger monitor retrieves the message. The trigger monitor uses the information in the trigger message. It issues a command to start the application that is to retrieve the messages arriving on the application queue, passing it information contained in the trigger message header, which includes the name of the application queue.

On all platforms, a special trigger monitor known as the channel initiator is responsible for starting channels.

z/OS On z/OS, the channel initiator is typically started manually, or it can be done automatically when a queue manager starts by changing CSQINP2 in the queue manager startup JCL.

Multi On Multiplatforms, the channel initiator is started automatically when the queue manager starts, or it can be started manually with the `runmqchi` command.

For more information, see “Trigger monitors” on page 248.

To understand how triggering works, consider Figure 48, which is an example of trigger type FIRST (MQTT_FIRST).

In Figure 48, the sequence of events is:




Figure 48. Flow of application and trigger messages

1. Application A, which can be either local or remote to the queue manager, puts a message on the application queue. No application has this queue open for input. However, this fact is relevant only to trigger type FIRST and DEPTH.

2. The queue manager checks to see if the conditions are met under which it has to generate a trigger event. They are, and a trigger event is generated. Information held within the associated process definition object is used when creating the trigger message.
3. The queue manager creates a trigger message and puts it on the initiation queue associated with this application queue, but only if an application (trigger monitor) has the initiation queue open for input.
4. The trigger monitor retrieves the trigger message from the initiation queue.
5. The trigger monitor issues a command to start application B (the server application).
6. Application B opens the application queue and retrieves the message.

Note:

1. If the application queue is open for input, by any program, and has triggering set for FIRST or DEPTH, no trigger event occurs because the queue is already being served.
2. If the initiation queue is not open for input, the queue manager does not generate any trigger messages; it waits until an application opens the initiation queue for input.
3. When using triggering for channels, use trigger type FIRST or DEPTH.
4. Triggered applications run under the user ID and group of the user who started the trigger monitor, the CICS user, or the user who started the queue manager.

So far, the relationship between the queues within triggering has been only on a one to one basis. Consider Figure 49 on page 238.




Figure 49. Relationship of queues within triggering

An application queue has a process definition object associated with it that holds details of the application that will process the message. The queue manager places the information in the trigger message, so only one initiation queue is necessary. The trigger monitor extracts this information from the trigger message and starts the relevant application to deal with the message on each application queue.

Remember that, if you want to trigger the start of a channel, you do not need to define a process definition object. The transmission queue definition can determine the channel to be triggered.

Use the following links to find out more about starting IBM MQ applications using triggers:

- “Prerequisites for triggering” on page 239
- “Conditions for a trigger event” on page 241
- “Controlling trigger events” on page 245
- “Designing an application that uses triggered queues” on page 247
- “Trigger monitors” on page 248

- “Properties of trigger messages” on page 252
- “When triggering does not work” on page 253

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222


This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

 “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

 “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

This information helps you to write IMS applications using IBM MQ.

Prerequisites for triggering:

Use this information to learn about the steps to take before using triggering.

Before your application can take advantage of triggering, complete the following steps:

1. Either:

- a. Create an initiation queue for your application queue. For example:

```
DEFINE QLOCAL (initiation.queue) REPLACE +
      LIKE (SYSTEM.DEFAULT.INITIATION.QUEUE) +
      DESCR ('initiation queue description')
```

or

- b. Determine the name of a local queue that exists and can be used by your application (typically, this name is SYSTEM.DEFAULT.INITIATION.QUEUE or, if you are starting channels with triggers, SYSTEM.CHANNEL.INITQ), and specify its name in the *InitiationQName* field of the application queue.

2. Associate the initiation queue with the application queue. A queue manager can own more than one initiation queue. You might want some of your application queues to be served by different programs, in which case, you can use one initiation queue for each serving program, although you do not have to. Here is an example of how to create an application queue:

```
DEFINE QLOCAL (application.queue) REPLACE +
      LIKE (SYSTEM.DEFAULT.LOCAL.QUEUE) +
      DESCR ('appl queue description') +
```

```

INITQ ('initiation.queue')      +
PROCESS ('process.name')       +
TRIGGER                          +
TRIGTYPE (FIRST)

```

IBM i Here is an extract from a CL program for IBM MQ for IBM i that creates an initiation queue:

```

/* Queue used by AMQSINQA          */
  CRTMQMQ      QNAME('SYSTEM.SAMPLE.INQ')  +
               QTYPE(*LCL) REPLACE(*YES)   +
               MQMNAME                      +
               TEXT('queue for AMQSINQA')   +
               SHARE(*YES) /* Shareable */+
               DFTMSGPST(*YES)/* Persistent messages OK */+
               +
               TRGENBL(*YES) /* Trigger control on */+
               TRGTYPE(*FIRST)/* Trigger on first message*/+
               PRCNAME('SYSTEM.SAMPLE.INQPROCESS')  +
               INITQNAME('SYSTEM.SAMPLE.TRIGGER')

```

- If you are triggering an application, create a process definition object to contain information relating to the application that is to serve your application queue. For example, to trigger-start a CICS payroll transaction called PAYR:

```

DEFINE PROCESS (process.name) +
  REPLACE +
  DESCR ('process description') +
  APPLICID ('PAYR') +
  APPLTYPE (CICS) +
  USERDATA ('Payroll data')

```

IBM i Here is an extract from a CL program for IBM MQ for IBM i that creates a process definition object:

```



/* Process definition          */
  CRTMQMPRC   PRCNAME('SYSTEM.SAMPLE.INQPROCESS')  +
               REPLACE(*YES)                      +
               MQMNAME                              +
               TEXT('trigger process for AMQSINQA') +
               ENVDATA('JOBPTY(3)') /* Submit parameter */+
               APPID('AMQSINQA') /* Program name */


```

When the queue manager creates a trigger message, it copies information from the attributes of the process definition object into the trigger message.

Platform	To create a process definition object
UNIX, Linux, and Windows systems	Use DEFINE PROCESS or use SYSTEM.DEFAULT.PROCESS and modify using ALTER PROCESS
z/OS z/OS	Use DEFINE PROCESS (see sample code in step 3), or use the operations and control panels.
IBM i IBM i	Use a CL program containing code as in step 3.

- Optional: Create a transmission queue definition and use blanks for the **ProcessName** attribute. The **TrigData** attribute can contain the name of the channel to be triggered or it can be left blank. Except on IBM MQ for z/OS, if it is left blank, the channel initiator searches the channel definition files until it finds a channel that is associated with the named transmission queue. When the queue manager creates a trigger message, it copies information from the **TrigData** attribute of the transmission queue definition into the trigger message.
- If you have created a process definition object to specify properties of the application that is to serve your application queue, associate the process object with your application queue by naming it in the **ProcessName** attribute of the queue.

Platform	Use commands
UNIX, Linux, and Windows systems	ALTER QLOCAL
 z/OS	ALTER QLOCAL
 IBM i	CHGMQMQ

6. Start instances of the trigger monitors  (or trigger servers in IBM MQ for IBM i) that are to serve the initiation queues you have defined. See “Trigger monitors” on page 248 for more information.

If you want to be aware of any undelivered trigger messages, make sure that your queue manager has a dead-letter (undelivered-message) queue defined. Specify the name of the queue in the *DeadLetterQName* queue manager field.

You can then set the trigger conditions that you require, using the attributes of the queue object that defines your application queue. For more information, see “Controlling trigger events” on page 245.

Conditions for a trigger event:

The queue manager creates a trigger message when the conditions detailed in this topic are satisfied.

References to shared queues in this topic mean shared queues in a queue-sharing group, only available on IBM MQ for z/OS.

The following conditions cause the queue manager to create a trigger message:

1. A message is *put* on a queue.
2. The message has a priority greater than or equal to the threshold trigger priority of the queue. This priority is set in the **TriggerMsgPriority** local queue attribute; if it is set to zero, any message qualifies.
3. The number of messages on the queue with priority greater than or equal to *TriggerMsgPriority* was previously, depending on *TriggerType*:
 - Zero (for trigger type MQTT_FIRST)
 - Any number (for trigger type MQTT EVERY)
 - *TriggerDepth* minus 1 (for trigger type MQTT_DEPTH)

Note:

- a. For non-shared local queues, the queue manager counts both committed and uncommitted messages when it assesses whether the conditions for a trigger event exist. Consequently an application might be started when there are no messages for it to retrieve because the messages on the queue have not been committed. In this situation, consider using the wait option with a suitable *WaitInterval*, so that the application waits for its messages to arrive.
 - b. For local shared queues, the queue manager counts committed messages only.
4. For triggering of type FIRST or DEPTH, no program has the application queue open for removing messages (that is, the **OpenInputCount** local queue attribute is zero).

Note:

- a. For shared queues, special conditions apply when multiple queue managers have trigger monitors running against a queue. In this situation, if one or more queue managers have the queue open for input shared, the trigger criteria on the other queue managers are treated as

TriggerType MQTT_FIRST and *TriggerMsgPriority* zero. When all the queue managers close the queue for input, the trigger conditions revert to those conditions specified in the queue definition.

An example scenario affected by this condition is multiple queue managers QM1, QM2, and QM3 with a trigger monitor running for an application queue A. A message arrives on A satisfying the conditions for triggering, and a trigger message is generated on the initiation queue. The trigger monitor on QM1 gets the trigger message and triggers an application. The triggered application opens the application queue for shared input. From this point on the trigger conditions for application queue A are evaluated as *TriggerType* MQTT_FIRST, and *TriggerMsgPriority* zero on queue managers QM2 and QM3, until QM1 closes the application queue.

- b. For shared queues, this condition is applied for each queue manager. That is, a queue manager's *OpenInputCount* for a queue must be zero for a trigger message to be generated for the queue by that queue manager. However, if any queue manager in the queue-sharing group has the queue open using the MQOO_INPUT_EXCLUSIVE option, no trigger message is generated for that queue by any of the queue managers in the queue-sharing group.

The change in how the trigger conditions are evaluated occurs when the triggered application opens the queue for input. In scenarios where there is only one trigger monitor running, other applications can have the same effect because they similarly open the application queue for input. It does not matter whether the application queue was opened by an application that is started by a trigger monitor, or by some other application; it is the fact that the queue is open for input on another queue manager that causes the change in trigger criteria.

5. On IBM MQ for z/OS, if the application queue is one with a **Usage** attribute of MQUS_NORMAL, get requests for it are not inhibited (that is, the **InhibitGet** queue attribute is MQQA_GET_ALLOWED). Also, if the triggered application queue is one with a **Usage** attribute of MQUS_XMITQ, get requests for it are not inhibited.
6. Either:
 - The **ProcessName** local queue attribute for the queue is not blank, and the process definition object identified by that attribute has been created, or
 - The **ProcessName** local queue attribute for the queue is all blank, but the queue is a transmission queue. As the process definition is optional, the **TriggerData** attribute might also contain the name of the channel to be started. In this case, the trigger message contains attributes with the following values:
 - **QName**: queue name
 - **ProcessName**: blanks
 - **TriggerData**: trigger data
 - **ApplType**: MQAT_UNKNOWN
 - **ApplId**: blanks
 - **EnvData**: blanks
 - **UserData**: blanks
7. An initiation queue has been created, and has been specified in the **InitiationQName** local queue attribute. Also:
 - Get requests are not inhibited for the initiation queue (that is, the value of the **InhibitGet** queue attribute is MQQA_GET_ALLOWED).
 - Put requests must not be inhibited for the initiation queue (that is, the value of the **InhibitPut** queue attribute must be MQQA_PUT_ALLOWED).
 - The value of the **Usage** attribute of the initiation queue must be MQUS_NORMAL.
 - In environments where dynamic queues are supported, the initiation queue must not be a dynamic queue that has been marked as logically deleted.
8. A trigger monitor currently has the initiation queue open for removing messages (that is, the **OpenInputCount** local queue attribute is greater than zero).

9. The trigger control (**TriggerControl** local queue attribute) for the application queue is set to MQTC_ON. To do this, set the **trigger** attribute when you define your queue, or use the ALTER QLOCAL command.
10. The trigger type (**TriggerType** local queue attribute) is not MQTT_NONE.
If all the required conditions are met, and the message that caused the trigger condition is put as part of a unit of work, the trigger message does not become available for retrieval by the trigger monitor application until the unit of work completes, whether the unit of work is committed or, for trigger type MQTT_FIRST or MQTT_DEPTH, backed out.
11. A suitable message is placed on the queue, for a **TriggerType** of MQTT_FIRST or MQTT_DEPTH, and the queue:
 - Was not previously empty (MQTT_FIRST), or
 - Had **TriggerDepth** or more messages (MQTT_DEPTH)

and conditions 2 on page 241 through 10 (excluding 3 on page 241) are satisfied, if in the case of MQTT_FIRST a sufficient interval (**TriggerInterval** queue manager attribute) has elapsed since the last trigger message was written for this queue.

This is to allow for a queue server that ends before processing all the messages on the queue. The purpose of the trigger interval is to reduce the number of duplicate trigger messages that are generated.

Note: If you stop and restart the queue manager, the **TriggerInterval** timer is reset. There is a small window during which it is possible to produce two trigger messages. The window exists when the trigger attribute of the queue is set to enabled at the same time as a message arrives and the queue was not previously empty (MQTT_FIRST) or had **TriggerDepth** or more messages (MQTT_DEPTH).

12. The only application serving a queue issues an MQCLOSE call, for a **TriggerType** of MQTT_FIRST or MQTT_DEPTH, and there is at least:
 - One (MQTT_FIRST), or
 - **TriggerDepth** (MQTT_DEPTH)

messages on the queue of sufficient priority (condition 2 on page 241), and conditions 6 on page 242 through 10 are also satisfied.

This is to allow for a queue server that issues an MQGET call, finds the queue empty, and so ends; however, in the interval between the MQGET and the MQCLOSE calls, one or more messages arrive.

Note:

- a. If the program serving the application queue does not retrieve all the messages, this can cause a closed loop. Each time that the program closes the queue, the queue manager creates another trigger message that causes the trigger monitor to start the server program again.
 - b. If the program serving the application queue backs out its get request (or if the program abends) before it closes the queue, the same happens. However, if the program closes the queue before backing out the get request, and the queue is otherwise empty, no trigger message is created.
 - c. To prevent such a loop occurring, use the *BackoutCount* field of MQMD to detect messages that are repeatedly backed out. For more information, see “Messages that are backed out” on page 39.
13. The following conditions are satisfied using MQSET or a command:
 - a.
 - **TriggerControl** is changed to MQTC_ON, or
 - **TriggerControl** is already MQTC_ON and the value of either **TriggerType**, **TriggerMsgPriority**, or **TriggerDepth** (if relevant) is changed,
 and there is at least:
 - One (MQTT_FIRST or MQTT EVERY), or
 - **TriggerDepth** (MQTT_DEPTH)

messages on the queue of sufficient priority (condition 2 on page 241), and conditions 4 on page 241 through 10 on page 243 (excluding 8 on page 242) are also satisfied.

This is to allow for an application or operator changing the triggering criteria, when the conditions for a trigger to occur are already satisfied.

- b. The value of the **InhibitPut** queue attribute of an initiation queue changes from MQQA_PUT_INHIBITED to MQQA_PUT_ALLOWED, and there is at least:

- One (MQTT_FIRST or MQTT_EVERY), or
- **TriggerDepth** (MQTT_DEPTH)

messages of sufficient priority (condition 2 on page 241) on any of the queues for which this is the initiation queue, and conditions 4 on page 241 through 10 on page 243 are also satisfied. (One trigger message is generated for each such queue satisfying the conditions.)

This is to allow for trigger messages not being generated because of the MQQA_PUT_INHIBITED condition on the initiation queue, but this condition now having been changed.

- c. The value of the **InhibitGet** queue attribute of an application queue changes from MQQA_GET_INHIBITED to MQQA_GET_ALLOWED, and there is at least:

- One (MQTT_FIRST or MQTT_EVERY), or
- **TriggerDepth** (MQTT_DEPTH)

messages of sufficient priority (condition 2 on page 241) on the queue, and conditions 4 on page 241 through 10 on page 243, excluding 5 on page 242, are also satisfied.

This allows applications to be triggered only when they can retrieve messages from the application queue.

- d. A trigger-monitor application issues an MQOPEN call for input from an initiation queue, and there is at least:

- One (MQTT_FIRST or MQTT_EVERY), or
- **TriggerDepth** (MQTT_DEPTH)

messages of sufficient priority (condition 2 on page 241) on any of the application queues for which this is the initiation queue, and conditions 4 on page 241 through 10 on page 243 (excluding 8 on page 242) are also satisfied, and no other application has the initiation queue open for input (one trigger message is generated for each such queue satisfying the conditions).

This is to allow for messages arriving on queues while the trigger monitor is not running, and for the queue manager restarting and trigger messages (which are nonpersistent) being lost.

14. MSGDLVSQ is set correctly. If you set MSGDLVSQ=FIFO, messages are delivered to the queue in a First In First Out basis. The priority of the message is ignored and the default priority of the queue is assigned to the message. If **TriggerMsgPriority** is set to a higher value than the default priority of the queue, no messages are triggered. If **TriggerMsgPriority** is set equal to or lower than the default priority of the queue, triggering occurs for type FIRST, EVERY, and DEPTH. For information about these types, see the description of the **TriggerType** field under “Controlling trigger events” on page 245.

If you set MSGDLVSQ=PRIORITY and the message priority is equal to or greater than the *TriggerMsgPriority* field, messages only count towards a trigger event. In this case, triggering occurs for type FIRST, EVERY, and DEPTH. As an example, if you put 100 messages of lower priority than the **TriggerMsgPriority**, the effective queue depth for triggering purposes is still zero. If you then put another message on the queue, but this time the priority is greater than or equal to the **TriggerMsgPriority**, the effective queue depth increases from zero to one and the condition for **TriggerType** FIRST is satisfied.

Note:

1. From step 12 on page 243 (where trigger messages are generated as a result of some event other than a message arriving on the application queue), the trigger message is not put as part of a unit of work. Also, if the **TriggerType** is MQTT_EVERY, and if there are one or more messages on the application queue, only one trigger message is generated.
2. If IBM MQ segments a message during MQPUT, a trigger event will not be processed until all the segments have been successfully placed on the queue. However, once message segments are on the queue, IBM MQ treats them as individual messages for triggering purposes. For example, a single logical message split into three pieces causes only one trigger event to be processed when it is first MQPUT and segmented. However, each of the three segments causes their own trigger events to be processed as they are moved through the IBM MQ network.

Controlling trigger events:

You control trigger events using some of the attributes that define your application queue. This information also gives examples of using the trigger types: EVERY, FIRST, and DEPTH.

You can enable and disable triggering, and you can select the number or priority of the messages that count toward a trigger event. There is a full description of these attributes in Attributes of objects.

The relevant attributes are:

TriggerControl

Use this attribute to enable and disable triggering for an application queue.

TriggerMsgPriority

The minimum priority that a message must have for it to count toward a trigger event. If a message of priority less than *TriggerMsgPriority* arrives on the application queue, the queue manager ignores the message when it determines whether to create a trigger message. If *TriggerMsgPriority* is set to zero, all messages count toward a trigger event.

TriggerType

In addition to the trigger type NONE (which disables triggering just like setting the *TriggerControl* to OFF), you can use the following trigger types to set the sensitivity of a queue to trigger events:

EVERY

A trigger event occurs every time that a message arrives on the application queue. Use this type of trigger if you want multiple instances of an application started.

FIRST A trigger event occurs only when the number of messages on the application queue changes from zero to one. Use this type of trigger if you want a serving program to start when the first message arrives on a queue, continue until there are no more messages to process, then end. You must always process the queue until it is empty. See also "Special case of trigger type FIRST" on page 246.

DEPTH

A trigger event occurs only when the number of messages on the application queue reaches the value of the **TriggerDepth** attribute. A typical use of this type of triggering is to start a program when all the replies to a set of requests are received.

Triggering by depth: With triggering by depth, the queue manager disables triggering (using the *TriggerControl* attribute) after it creates a trigger message. Your application must re-enable triggering itself (by using the MQSET call) after this has happened.

The action of disabling triggering is not under syncpoint control, so triggering cannot be re-enabled by backing out a unit of work. If a program backs out a put request that caused a trigger event, or if the program abends, you must re-enable triggering by using the MQSET call or the ALTER QLOCAL command.

TriggerDepth

The number of messages on a queue that causes a trigger event when using triggering by depth.

The conditions that must be satisfied for a queue manager to create a trigger message are described in “Conditions for a trigger event” on page 241.

Example of the use of trigger type EVERY

Consider an application that generates requests for motor insurance. The application might send request messages to a number of insurance companies, specifying the same reply-to queue each time. It might set a trigger of type EVERY on this reply-to queue so that each time a reply arrives, the reply might trigger an instance of the server to process the reply.

Example of the use of trigger type FIRST

Consider an organization with a number of branch offices that each transmit details of the days business to the head office. They all do this at the same time, at the end of the working day, and at the head office there is an application that processes the details from all the branch offices. The first message to arrive at the head office could cause a trigger event that starts this application. This application would continue processing until there are no more messages on its queue.

Example of the use of trigger type DEPTH

Consider a travel agency application that creates a single request to confirm a flight reservation, to confirm a reservation for a hotel room, to rent a car, and to order some travelers checks. The application might separate these items into four request messages, sending each to a separate destination. It might set a trigger of type DEPTH on its reply-to queue (with the depth set to the value 4), so that it is restarted only when all four replies have arrived.

If another message (possibly from a different request) arrives on the reply-to queue before the last of the four replies, the requesting application is triggered early. To avoid this, when using DEPTH triggering to collect multiple replies to a request, always use a new reply-to queue for each request.

Special case of trigger type FIRST

With trigger type FIRST, if there is already a message on the application queue when another message arrives, the queue manager does not typically create another trigger message.

However, the application serving the queue might not actually open the queue (for example, the application might end, possibly because of a system problem). If an incorrect application name has been put into the process definition object, the application serving the queue will not pick up any of the messages. In these situations, if another message arrives on the application queue, there is no server running to process this message (and any other messages on the queue).

To deal with this, the queue manager creates further trigger messages under the following circumstances:

- If another message arrives on the application queue, but only if a predefined time interval has elapsed since the queue manager created the last trigger message for that queue. This time interval is defined in the queue manager attribute *TriggerInterval*. Its default value is 999 999 999 milliseconds.
- On IBM MQ for z/OS, application queues that name an open initiation queue are scanned periodically. If *TRIGINT* milliseconds have passed since the last trigger message was sent and the queue satisfies the conditions for a trigger event and *CURDEPTH* is greater than zero, a trigger message is generated. This process is called backstop triggering.

Consider the following points when deciding on a value for the trigger interval to use in your application:

- If you set *TriggerInterval* to a low value, and there is no application serving the application queue, trigger type FIRST might behave like trigger type EVERY. This depends on the rate that messages are being put onto the application queue, which in turn might depend on other system activity. This is because, if the trigger interval is very small, another trigger message is generated each time that a message is put onto the application queue, even though the trigger type is FIRST, not EVERY. (Trigger type FIRST with a trigger interval of zero is equivalent to trigger type EVERY.)
- On IBM MQ for z/OS if you set *TRIGINT* to a low value, and there is no application serving the trigger type FIRST application queue, backstop triggering will generate a trigger message each time the periodic scan of application queues that name open initiation queues takes place.
- If a unit of work is backed out (see Trigger messages and units of work) and the trigger interval has been set to a high value (or the default value), one trigger message is generated when the unit of work is backed out. However, if you have set the trigger interval to a low value or to zero (causing trigger type FIRST to behave like trigger type EVERY) many trigger messages can be generated. If the unit of work is backed out, all the trigger messages are still made available. The number of trigger messages that are generated depends on the trigger interval. If the trigger interval is set to zero, the maximum number of messages are generated.

Designing an application that uses triggered queues:

You have seen how to set up, and control, triggering for your applications. Here are some tips to consider when you design your application.

Trigger messages and units of work

Trigger messages created because of trigger events that are not part of a unit of work are put on the initiation queue, outside any unit of work, with no dependence on any other messages, and are available for retrieval by the trigger monitor immediately.

Trigger messages created because of trigger events that are part of a unit of work are made available on the initiation queue when the UOW is resolved, whether the unit of work is committed or backed out

If the queue manager fails to put a trigger message on an initiation queue, it will be put on the dead-letter (undelivered-message) queue.

Note:

1. The queue manager counts both committed and uncommitted messages when it assesses whether the conditions for a trigger event exist.

With triggering of type FIRST or DEPTH, trigger messages are made available even if the unit of work is backed out so that a trigger message is always available when the required conditions are met. For example, consider a put request within a unit of work for a queue that is triggered with trigger type FIRST. This causes the queue manager to create a trigger message. If another put request occurs, from another unit of work, this does not cause another trigger event because the number of messages on the application queue has now changed from one to two, which does not satisfy the conditions for a trigger event. Now if the first unit of work is backed out, but the second is committed, a trigger message is still created.

However, this means that trigger messages are sometimes created when the conditions for a trigger event are not satisfied. Applications that use triggering must always be prepared to handle this situation. It is recommended that you use the wait option with the MQGET call, setting the *WaitInterval* to a suitable value.

Created trigger messages are always made available, whether the unit of work is backed out or committed.

2. For local shared queues (that is, shared queues in a queue-sharing group) the queue manager counts committed messages only.

Getting messages from a triggered queue

When you design applications that use triggering, be aware that there might be a delay between a trigger monitor starting a program and other messages becoming available on the application queue. This can happen when the message that causes the trigger event is committed before the others.

To allow time for messages to arrive, always use the wait option when you use the MQGET call to remove messages from a queue for which trigger conditions are set. The *WaitInterval* must be sufficient to allow for the longest reasonable time between a message being put and that put call being committed. If the message is arriving from a remote queue manager, this time is affected by:

- The number of messages that are put before being committed
- The speed and availability of the communication link
- The sizes of the messages

For an example of a situation where you should use the MQGET call with the wait option, consider the same example that we used when describing units of work. This was a put request within a unit of work for a queue that is triggered with trigger type FIRST. This event causes the queue manager to create a trigger message. If another put request occurs, from another unit of work, this does not cause another trigger event because the number of messages on the application queue has not changed from zero to one. Now if the first unit of work is backed out, but the second is committed, a trigger message is still created. So the trigger message is created at the time that the first unit of work is backed out. If there is a significant delay before the second message is committed, the triggered application might need to wait for it.

With triggering of type DEPTH, a delay can occur even if all relevant messages are eventually committed. Suppose that the **TriggerDepth** queue attribute has the value 2. When two messages arrive on the queue, the second causes a trigger message to be created. However, if the second message is the first to be committed, it is at that time that the trigger message becomes available. The trigger monitor starts the server program, but the program can retrieve only the second message until the first one is committed. So the program might need to wait for the first message to be made available.

Design your application so that it terminates if no messages are available for retrieval when your wait interval expires. If one or more messages arrive later, rely on your application being retriggered to process them. This method prevents applications being idle, and unnecessarily using resources.

Trigger monitors:

To a queue manager, a trigger monitor is like any other application that serves a queue. However, a trigger monitor serves initiation queues.

A trigger monitor is usually a continuously-running program. When a trigger message arrives on an initiation queue, the trigger monitor retrieves that message. It uses information in the message to issue a command to start the application that is to process the messages on the application queue.

The trigger monitor must pass sufficient information to the program that it is starting so that the program can perform the correct actions on the correct application queue.

A channel initiator is an example of a special type of trigger monitor for message channel agents. In this situation however, you must use either trigger type FIRST or DEPTH.

Trigger monitors on UNIX and Windows systems:

UNIX

Windows

This topic contains information about trigger monitors provided on UNIX and Windows systems.

The following trigger monitors are provided for the server environment:

amqstrg0

This is a sample trigger monitor that provides a subset of the function provided by **runmqtrm**. See "Using the sample programs on Multiplatforms" on page 457 for more information about **amqstrg0**.

runmqtrm

The syntax of this command is **runmqtrm** [*-m QMgrName*] [*-q InitQ*], where *QMGrName* is the queue manager and *InitQ* is the initiation queue. The default queue is SYSTEM.DEFAULT.INITIATION.QUEUE on the default queue manager. It calls programs for the appropriate trigger messages. This trigger monitor supports the default application type.

The command string passed by the trigger monitor to the operating system is built as follows:

1. The *ApplId* from the relevant PROCESS definition (if created)
2. The MQTMC2 structure, enclosed in double quotation marks
3. The *EnvData* from the relevant PROCESS definition (if created)

where *ApplId* is the name of the program to run as it would be entered on the command line.

The parameter passed is the MQTMC2 character structure. A command string is invoked that has this string, exactly as provided, in double quotation marks, in order that the system command will accept it as one parameter.

The trigger monitor does not look to see if there is another message on the initiation queue until the completion of the application that it has just started. If the application has much processing to do, the trigger monitor might not be able to keep up with the number of trigger messages arriving. You have two options:

- Have more trigger monitors running
- Run the started applications in the background

If you have more trigger monitors running, you can control the maximum number of applications that can run at any one time. If you run applications in the background, there is no restriction imposed by IBM MQ on the number of applications that can run.

To run the started application in the background on Windows systems, within the *ApplId* field, prefix the name of your application with a START command. For example:

```
START ?B AMQSECHA
```

UNIX

To run the started application in the background on UNIX, put an & at the end of the *EnvData* of the PROCESS definition.

Note: Windows Where a Windows path has spaces as a part of the path name, these should be enclosed in quotation marks (") to ensure that it is handled as a single argument. For example, "C:\Program Files\Application Directory\Application.exe".

The following is an example of an APPLICID string where the file name includes spaces as a part of the path:

```
START "" /B "C:\Program Files\Application Directory\Application.exe"
```

The syntax of the Windows START command in the example includes an empty string enclosed in double quotation marks. START specifies that the first argument in quotation marks will be

treated as the title of the new command. To ensure that Windows does not mistake the application path for a 'title' argument, add a title string enclosed in double quotation marks to the command before the application name.

The following trigger monitors are provided for the IBM MQ client:

runmqtmc

This is the same as runmqtrm except that it links with the IBM MQ MQI client libraries.

For CICS:

The amqltmc0 trigger monitor is provided for CICS. It works in the same way as the standard trigger monitor, runmqtrm, but you run it in a different way and it triggers CICS transactions.

This topic applies only to Windows, UNIX, and Linux x86-64 systems.

It is supplied as a CICS program; define it with a 4-character transaction name. Enter the 4-character name to start the trigger monitor. It uses the default queue manager (as named in the qm.ini file or, on IBM MQ for Windows, the registry), and the SYSTEM.CICS.INITIATION.QUEUE.

If you want to use a different queue manager or queue, build the trigger monitor MQTMC2 structure: this requires you to write a program using the EXEC CICS START call, because the structure is too long to add as a parameter. Then, pass the MQTMC2 structure as data to the START request for the trigger monitor.

When you use the MQTMC2 structure, you need to supply only the *StrucId*, *Version*, *QName*, and **QMgrName** parameters to the trigger monitor as it does not reference any other fields.


Messages are read from the initiation queue and used to start CICS transactions, using EXEC CICS START, assuming the APPL_TYPE in the trigger message is MQAT_CICS. The reading of messages from the initiation queue is performed under CICS syncpoint control.

Messages are generated when the monitor starts and stops, and when an error occurs. These messages are sent to the CSMT transient data queue.

Here are the available versions of the trigger monitor:

Version	Use
amqltmc0	TXSeries for AIX, HP-UX, Linux x86-64 systems, and Oracle Solaris Version 5.1
amqltmc4	TXSeries for Windows, Version 5.1
amqltmcc	Client bound version of the CICS trigger monitor


If you need a trigger monitor for other environments, write a program that can process the trigger messages that the queue manager puts on the initiation queues. Such a program should perform the following actions:

1. Use the MQGET call to wait for a message to arrive on the initiation queue.
2. Examine the fields of the MQTM structure of the trigger message to find the name of the application to start, and the environment in which it runs.
3. Issue an environment-specific start command.  For example, in z/OS batch, submit a job to the internal reader.
4. Convert the MQTM structure to the MQTMC2 structure if required.
5. Pass either the MQTMC2 or MQTM structure to the started application. This can contain user data.

6. Associate with your application queue the application that is to serve that queue. You do this by naming the process definition object (if created) in the **ProcessName** attribute of the queue.

 Use DEFINE QLOCAL or ALTER QLOCAL. On IBM i, you can also use CRTMQMQ or CHGMQMQ.

For more information about the trigger monitor interface, see MQTMC2.

Trigger monitors on IBM i: 

On IBM i, instead of the **runmqtrm** control command, use the IBM MQ for IBM i CL command **STRMQMTRM**.

Use the STRMQMTRM command as follows:

```
STRMQMTRM INITQNAME(InitQ) MQMNAME(QMgrName)
```

Details are as for runmqtrm.

The following sample programs are also provided, which you can use as models to write your own trigger monitors:

AMQSTRG4

This is a trigger monitor that submits an IBM i job for the process that is to be started, but this means that there is additional processing associated with each trigger message.

AMQSERV4

This is a trigger server. For each trigger message, this server runs the command for the process in its own job, and can call CICS transactions.

Both the trigger monitor and the trigger server pass an MQTMC2 structure to the programs that they start. For a description of this structure, see MQTMC2. Both of these samples are delivered in both source and executable forms.

Because these trigger monitors can invoke only native IBM i programs, they cannot trigger Java programs directly, because Java classes are located in the IFS. However, Java programs can be triggered indirectly by triggering a CL program that then invokes the Java program and passes across the TMC2 structure. The minimum size of the TMC2 structure is 732 bytes.

Here is the source of a sample CLP:

```
PGM PARM(&TMC2)
DCL &TMC2 *CHAR LEN(800)
ADDENVVAR ENVVAR(TM) VALUE(&TMC2)
QSH CMD('java_pgmname $TM')
RMVENVVAR ENVVAR(TM)
ENDPGM
```

The following trigger monitor program is provided for the IBM MQ MQI client: RUNMQTMC

Call the RUNMQTMC as follows:

```
CALL PGM(QMQM/RUNMQTMC) PARM('-m' QMgrName '-q' InitQ)
```

Properties of trigger messages:

The following topics describe some other properties of trigger messages.

- “Persistence and priority of trigger messages”
- “Queue manager restart and trigger messages”
- “Trigger messages and changes to object attributes”
- “Format of trigger messages”

Persistence and priority of trigger messages

Trigger messages are not persistent because there is no requirement for them to be so.

However, the conditions for generating triggering events do persist, so trigger messages are generated whenever these conditions are met. If a trigger message is lost, the continued existence of the application message on the application queue guarantees that the queue manager generates a trigger message as soon as all the conditions are met.

If a unit of work is rolled back, any trigger messages it generated are always delivered.

Trigger messages take the default priority of the initiation queue.

Queue manager restart and trigger messages

Following the restart of a queue manager, when an initiation queue is next opened for input, a trigger message can be put to this initiation queue if an application queue associated with it has messages on it, and is defined for triggering.

Trigger messages and changes to object attributes

Trigger messages are created according to the values of the trigger attributes in force at the time of the trigger event.

If the trigger message is not made available to a trigger monitor until later (because the message that caused it to be generated was put within a unit of work), any changes to the trigger attributes in the meantime have no effect on the trigger message. In particular, disabling triggering does not prevent a trigger message being made available once it has been created. Also, the application queue might no longer exist at the time that the trigger message is made available.

Format of trigger messages

The format of a trigger message is defined by the MQTM structure.

This has the following fields, which the queue manager fills when it creates the trigger message, using information in the object definitions of the application queue and of the process associated with that queue:

StrucId

The structure identifier.

Version

The version of the structure.

QName The name of the application queue on which the trigger event occurred. When the queue manager creates a trigger message, it fills this field using the **QName** attribute of the application queue.

ProcessName

The name of the process definition object that is associated with the application queue. When the queue manager creates a trigger message, it fills this field using the **ProcessName** attribute of the application queue.

TriggerData

A free-format field for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the **TriggerData** attribute of the application queue. On any IBM MQ product except IBM MQ for z/OS, this field can be used to specify the name of the channel to be triggered.

ApplType

The type of the application that the trigger monitor is to start. When the queue manager creates a trigger message, it fills this field using the **ApplType** attribute of the process definition object identified in *ProcessName*.

ApplId A character string that identifies the application that the trigger monitor is to start. When the queue manager creates a trigger message, it fills this field using the **ApplId** attribute of the process definition object identified in *ProcessName*. When you use trigger monitor CKTI or CSQQTRMN supplied by IBM MQ for z/OS, the **ApplId** attribute of the process definition object is a CICS or IMS transaction identifier.

EnvData

A character field containing environment-related data for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the **EnvData** attribute of the process definition object identified in *ProcessName*. The IBM MQ for z/OS-supplied trigger monitors (CKTI or CSQQTRMN) do not use this field, but other trigger monitors might choose to use it.

UserData


A character field containing user data for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the **UserData** attribute of the process definition object identified in *ProcessName*. This field can be used to specify the name of the channel to be triggered.

There is a full description of the trigger message structure in MQTM.

When triggering does not work:

A program is not triggered if the trigger monitor cannot start the program or the queue manager cannot deliver the trigger message. For example, the applid in the process object must specify that the program is to be started in the background; otherwise, the trigger monitor cannot start the program.

If a trigger message is created but cannot be put on the initiation queue (for example, because the queue is full or the length of the trigger message is greater than the maximum message length specified for the initiation queue), the trigger message is put instead on the dead-letter (undelivered message) queue.

If the put operation to the dead-letter queue cannot complete successfully, the trigger message is discarded and a warning message is sent  to the z/OS console or to the system operator or is put on the error log.

Putting the trigger message on the dead-letter queue might generate a trigger message for that queue. This second trigger message is discarded if it adds a message to the dead-letter queue.

If the program is triggered successfully but abends before it receives the message from the queue, use a trace utility (for example, CICS AUXTRACE if the program is running under CICS) to find the cause of the failure.

Working with the MQI and clusters

There are special options on calls and return codes that relate to clustering.

Use the following links to find out more about the options available on the calls and return codes for use with clusters:

- “MQOPEN and clusters”
- “MQPUT, MQPUT1 and clusters” on page 256
- “MQINQ and clusters” on page 256
- “MQSET and clusters” on page 257
- “Return codes” on page 257

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

▶ **z/OS** “Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

▶ **z/OS** “IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

This information helps you to write IMS applications using IBM MQ.

MQOPEN and clusters:

The queue to which a message is put to, or read from, when a cluster queue is opened, depends on the MQOPEN call.

Selecting the target queue

If you do not provide a queue manager name in the object descriptor, MQOD, the queue manager selects the queue manager to send the message to. If you do provide a queue manager name in the object descriptor, then messages are always sent to the queue manager you have selected.

If the queue manager is selecting the target queue manager, the selection depends on the binding options, MQOO_BIND_* and if a local queue exists. If there is a local instance of the queue, it is always opened in preference to a remote instance, unless the CLWLUSEQ attribute is set to ANY. Otherwise, the selection

depends on the binding options. 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.

If the queue manager is selecting the target queue manager, it does so in a round-robin fashion, using the workload management algorithm; see *Workload balancing in clusters*.

When the workload balancing algorithm is used depends on the way the cluster queue is opened:

- `MQOO_BIND_ON_OPEN` - the algorithm is used once at the time the queue is opened by the application.
- `MQOO_BIND_NOT_FIXED` - the algorithm is used for every message put to the queue.
- `MQOO_BIND_ON_GROUP` - the algorithm is used once at the start of each message group.

MQOO_BIND_ON_OPEN

The `MQOO_BIND_ON_OPEN` option on the `MQOPEN` call specifies that the target queue manager is to be fixed. Use the `MQOO_BIND_ON_OPEN` option if there are multiple instances of the same queue within a cluster. All messages put to the queue specifying the object handle returned from the `MQOPEN` call are directed to the same queue manager.

- Use the `MQOO_BIND_ON_OPEN` option if messages have affinities. For example, if a batch of messages is all to be processed by the same queue manager, specify `MQOO_BIND_ON_OPEN` when you open the queue. IBM MQ fixes the queue manager and the route to be taken by all messages put to that queue.
- If `MQOO_BIND_ON_OPEN` option is specified, the queue must be reopened for a new instance of the queue to be selected.

MQOO_BIND_NOT_FIXED

The `MQOO_BIND_NOT_FIXED` option on the `MQOPEN` call specifies that the target queue manager is not fixed. Messages written to the queue specifying the object handle returned from the `MQOPEN` call are routed to a queue manager at `MQPUT` time on a message-by-message basis. Use the `MQOO_BIND_NOT_FIXED` option if you do not want to force all your messages to be written to the same destination.

- Do not specify `MQOO_BIND_NOT_FIXED` and `MQMF_SEGMENTATION_ALLOWED` at the same time. If you do, the segments of your message might be delivered to different queue managers, scattered throughout the cluster.

MQOO_BIND_ON_GROUP

Allows an application to request that a group of messages is allocated to the same destination instance. This option is valid only for queues, and affects only cluster queues. If specified for a queue that is not a cluster queue, the option is ignored.

- Groups are only routed to a single destination when `MQPMO_LOGICAL_ORDER` is specified on the `MQPUT`. When `MQOO_BIND_ON_GROUP` is specified, but a message is not part of a logical group, `BIND_NOT_FIXED` behavior is used instead.

MQOO_BIND_AS_Q_DEF

If you do not specify either `MQOO_BIND_ON_OPEN`, `MQOO_BIND_NOT_FIXED` or `MQOO_BIND_ON_GROUP`, the default option is `MQOO_BIND_AS_Q_DEF`. Using `MQOO_BIND_AS_Q_DEF` causes the binding that is used for the queue handle to be taken from the `DefBind` queue attribute.

Relevance of MQOPEN options

The `MQOPEN` options `MQOO_BROWSE`, `MQOO_INPUT_*`, or `MQOO_SET` require a local instance of the cluster queue for `MQOPEN` to succeed.

The `MQOPEN` options `MQOO_OUTPUT`, `MQOO_BIND_*`, or `MQOO_INQUIRE` do not require a local instance of the cluster queue to succeed.

Resolved queue manager name

When a queue manager name is resolved at MQOPEN time, the resolved name is returned to the application. If the application tries to use this name on a subsequent MQOPEN call, it might find that it is not authorized to access the name.

MQPUT, MQPUT1 and clusters:

If MQ00_BIND_NOT_FIXED is specified on an MQOPEN the workload management routines chooses which destination MQPUT or MQPUT1 select.

If MQ00_BIND_NOT_FIXED is specified on an MQOPEN call, each subsequent MQPUT call invokes the workload management routine to determine which queue manager to send the message to. The destination and route to be taken are selected on a message-by-message basis. The destination and route might change after the message has been put if conditions in the network change. The MQPUT1 call always operates as though MQ00_BIND_NOT_FIXED were in effect, that is, it always invokes the workload management routine.

When the workload management routine has selected a queue manager, the local queue manager completes the put operation. The message can be placed on different queues:

1. If the destination is the local instance of the queue, the message is placed on the local queue.
2. If the destination is a queue manager in a cluster, the message is placed on a cluster transmission queue.
3. If the destination is a queue manager outside a cluster, the message is placed on a transmission queue with the same name as the target queue manager.

If MQ00_BIND_ON_OPEN is specified on the MQOPEN call, MQPUT calls do not invoke the workload management routine because the destination and route have already been selected.

MQINQ and clusters:

Which cluster queue is inquired upon depends upon the options you combine with MQ00_INQUIRE.

Before you can inquire on a queue, open it using the MQOPEN call and specify MQ00_INQUIRE.

To inquire on a cluster queue, use the MQOPEN call and combine other options with MQ00_INQUIRE. The attributes that can be inquired depend on whether there is a local instance of the cluster queue, and on how the queue is opened:

- Combining MQ00_BROWSE, MQ00_INPUT_*, or MQ00_SET with MQ00_INQUIRE requires a local instance of the cluster queue for the open to succeed. In this case you can inquire on all the attributes that are valid for local queues.
- Combining MQ00_OUTPUT with MQ00_INQUIRE, and specifying none of the preceding options, the instance opened is either:
 - The instance on the local queue manager, if there is one. In this case you can inquire on all the attributes that are valid for local queues.
 - An instance elsewhere in the cluster, if there is no local queue manager instance. In this case only the following attributes can be inquired on. The QType attribute has the value MQQT_CLUSTER in this case.
 - DefBind
 - DefPersistence
 - DefPriority
 - InhibitPut
 - QDesc

- QName
- QType

To inquire on the DefBind attribute of a cluster queue, use the MQINQ call with the selector MQIA_DEF_BIND. The value returned is either MQBND_BIND_ON_OPEN or MQBND_BIND_NOT_FIXED, or MQBND_BIND_ON_GROUP. Either MQBND_BIND_ON_OPEN or MQBND_BIND_ON_GROUP must be specified when using groups with clusters.

To inquire on the CLUSTER and CLUSNL attributes of the local instance of a queue, use the MQINQ call with the selector MQCA_CLUSTER_NAME or the selector MQCA_CLUSTER_NAMELIST.

Note: If you open a cluster queue without fixing the queue that MQOPEN has bound to, successive MQINQ calls might inquire on different instances of the cluster queue.

Related concepts:

“MQOPEN option for cluster queue” on page 106

The binding used for the queue handle is taken from the **DefBind** queue attribute, which can take the value MQBND_BIND_ON_OPEN, MQBND_BIND_NOT_FIXED, or MQBND_BIND_ON_GROUP.

MQSET and clusters:

The MQOPEN option MQ00_SET option requires there to be a local instance of a cluster queue for MQSET to succeed.

You cannot use the MQSET call to set the attributes of a queue elsewhere in the cluster.

You can open a local alias or remote queue defined with the cluster attribute and use the MQSET call. You can set the attributes of the local alias or remote queue. It does not matter if the target queue is a cluster queue defined on a different queue manager.

Return codes:

Return codes specific to clusters

MQRC_CLUSTER_EXIT_ERROR (2266 X'8DA')

An MQOPEN, MQPUT, or MQPUT1 call is issued to open a cluster queue or put a message on it. The cluster workload exit, defined by the ClusterWorkloadExit attribute of a queue manager, fails unexpectedly or does not respond in time.


A message is written to the system log on IBM MQ for z/OS giving more information about this error.

Subsequent MQOPEN, MQPUT, and MQPUT1 calls for this queue handle are processed as though the ClusterWorkloadExit attribute were blank.

MQRC_CLUSTER_EXIT_LOAD_ERROR (2267 X'8DB')

On z/OS, the cluster workload exit cannot be loaded.

A message is written to the system log and processing continues as though the ClusterWorkloadExit attribute is blank.

 On Multiplatforms, an MQCONN or MQCONNX call is issued to connect to a queue manager. The call fails because the cluster workload exit, defined by the queue manager ClusterWorkloadExit attribute of the queue manager, cannot be loaded.

MQRC_CLUSTER_PUT_INHIBITED (2268 X'8DC')

An MQOPEN call with the MQ00_OUTPUT and MQ00_BIND_ON_OPEN options in effect is issued for a cluster queue. All the instances of the queue in the cluster are currently put-inhibited by having

the `InhibitPut` attribute set to `MQQA_PUT_INHIBITED`. Because there are no queue instances available to receive messages, the `MQOPEN` call fails.

This reason code occurs only when both of the following statements are true:

- There is no local instance of the queue. If there is a local instance, the `MQOPEN` call succeeds, even if the local instance is put-inhibited.
- There is no cluster workload exit for the queue, or there is a cluster workload exit but it does not choose a queue instance. (If the cluster workload exit chooses a queue instance, the `MQOPEN` call succeeds, even if that instance is put-inhibited.)

If the `MQ00_BIND_NOT_FIXED` option is specified on the `MQOPEN` call, the call can succeed even if all the queues in the cluster are put-inhibited. However, a subsequent `MQPUT` call might fail if all the queues are still put-inhibited at the time of that call.

MQRC_CLUSTER_RESOLUTION_ERROR (2189 X'88D')

1. An `MQOPEN`, `MQPUT`, or `MQPUT1` call is issued to open a cluster queue or put a message on it. The queue definition cannot be resolved correctly because a response is required from the full repository queue manager but none is available.
2. An `MQOPEN`, `MQPUT`, `MQPUT1` or `MQSUB` call is issued for a topic object specifying `PUBSCOPE (ALL)` or `SUBSCOPE (ALL)`. The cluster topic definition cannot be resolved correctly because a response is required from the full repository queue manager but none is available.

MQRC_CLUSTER_RESOURCE_ERROR (2269 X'8DD')

An `MQOPEN`, `MQPUT`, or `MQPUT1` call is issued for a cluster queue. An error occurs while trying to use a resource required for clustering.

MQRC_NO_DESTINATIONS_AVAILABLE (2270 X'8DE')

An `MQPUT` or `MQPUT1` call is issued to put a message on a cluster queue. At the time of the call, there are no longer any instances of the queue in the cluster. The `MQPUT` fails and the message is not sent.

The error can occur if `MQ00_BIND_NOT_FIXED` is specified on the `MQOPEN` call that opens the queue, or `MQPUT1` is used to put the message.

MQRC_STOPPED_BY_CLUSTER_EXIT (2188 X'88C')

An `MQOPEN`, `MQPUT`, or `MQPUT1` call is issued to open or put a message on a cluster queue. The cluster workload exit rejects the call.

Using and writing applications on IBM MQ for z/OS

z/OS

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

This information explains the IBM MQ facilities available to programs running in each of the supported environments. In addition,

- For information about using the IBM MQ-CICS bridge, see [Using IBM MQ with CICS](#).
- For information about using IMS and the IMS bridge, see [“IMS and IMS bridge applications on IBM MQ for z/OS”](#) on page 57.

Use the following links to find out more about using and writing applications on IBM MQ for z/OS:

- [“Environment-dependent IBM MQ for z/OS functions”](#) on page 259
- [“Debugging facilities, syncpoint support, and recovery support”](#) on page 260
- [“The IBM MQ for z/OS interface with the application environment”](#) on page 261
- [“Writing z/OS UNIX System Services applications”](#) on page 262

- “Application programming with shared queues” on page 266

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

“IMS and IMS bridge applications on IBM MQ for z/OS” on page 57

This information helps you to write IMS applications using IBM MQ.

Environment-dependent IBM MQ for z/OS functions:

Use this information when considering IBM MQ for z/OS functions.

The main differences to be considered between IBM MQ functions in the environments in which IBM MQ for z/OS runs are:

- IBM MQ for z/OS supplies the following trigger monitors:
 - CKTI for use in the CICS environment
 - CSQQTRMN for use in the IMS environment

You must write your own module to start applications in other environments.

- Syncpointing using two-phase commit is supported in the CICS and IMS environments. It is also supported in the z/OS batch environment using transaction management and recoverable resource manager services (RRS). Single-phase commit is supported in the z/OS environment by IBM MQ itself.
- For the batch and IMS environments, the MQI provides calls to connect programs to, and to disconnect them from, a queue manager. Programs can connect to more than one queue manager.
- A CICS system can connect to only one queue manager. This can be made to happen when CICS is initiated if the subsystem name is defined in the CICS system startup job. The MQI connect and disconnect calls are tolerated, but have no effect, in the CICS environment.
- The API-crossing exit allows a program to intervene in the processing of all MQI calls. This exit is available in the CICS environment only.
- In CICS on multiprocessor systems, some performance advantage is gained because MQI calls can be executed under multiple z/OS TCBS. For more information, see the *Planning on z/OS IBM MQ for z/OS Concepts and Planning Guide*.

These features are summarized in Table 31 on page 260.

Table 31. z/OS environmental features

	CICS	IMS	Batch/TSO
Trigger monitor supplied	Yes	Yes	No
Two-phase commit	Yes	Yes	Yes
Single-phase commit	Yes	No	Yes
Connect/disconnect MQI calls	Tolerated	Yes	Yes
API-crossing exit	Yes	No	No

Note: Two-phase commit is supported in the Batch/TSO environment using RRS.

Debugging facilities, syncpoint support, and recovery support:

Use this information to learn about program debugging facilities, syncpoint support, and recovery support.

Program debugging facilities

IBM MQ for z/OS provides a trace facility that you can use to debug your programs in all environments.

Additionally, in the CICS environment you can use:

- The CICS Execution Diagnostic Facility (CEDF)
- The CICS Trace Control Transaction (CETR)
- The IBM MQ for z/OS API-crossing exit

On the z/OS platform, you can use any available interactive debugging tool that is supported by the programming language that you are using.

Syncpoint support

Synchronizing the start and end of units of work is necessary in a transaction processing environment so that transaction processing can be used safely.

This is fully supported by IBM MQ for z/OS in the CICS and IMS environments. Full support means cooperation between resource managers so that units of work can be committed or backed out in unison, under control of CICS or IMS. Examples of resource managers are Db2, CICS File Control, IMS, and IBM MQ for z/OS.

z/OS batch applications can use IBM MQ for z/OS calls to give a single-phase commit facility. This means that an application-defined set of queue operations can be committed, or backed out, without reference to other resource managers.

Two-phase commit is also supported in the z/OS batch environment using transaction management and recoverable resource manager services (RRS). For further information see Syncpoints in z/OS batch applications.

Recovery support

If the connection between a queue manager and a CICS or IMS system is broken during a transaction, some units of work might not be backed out successfully.

However, these units of work are resolved by the queue manager (under the control of the syncpoint manager) when its connection with the CICS or IMS system is reestablished.

The IBM MQ for z/OS interface with the application environment:

To allow applications running in different environments to send and receive messages through a message queuing network, IBM MQ for z/OS provides an *adapter* for each of the environments it supports.

These adapters are the interface between application programs and IBM MQ for z/OS subsystems. They allow the programs to use the MQI.

The batch adapter:

Use this information to learn about the batch adapter and the commit protocol it supports.

The *batch adapter* provides access to IBM MQ for z/OS resources for programs running in:

- Task (TCB) mode
- Problem or supervisor state
- Primary address space control mode

The programs must not be in cross-memory mode.

Connections between application programs and IBM MQ for z/OS are at the task level. The adapter provides a single connection thread from an application task control block (TCB) to IBM MQ for z/OS.

The adapter supports a single-phase commit protocol for changes made to resources owned by IBM MQ for z/OS ; it does not support multiphase-commit protocols.

The RRS batch adapter:

Use this information to learn about the RRS batch adapter and the two RRS batch adapters provided by IBM MQ.

The transaction management and recoverable resource manager services (RRS) adapter:

- Uses z/OS RRS for commit control.
- Supports simultaneous connections to multiple IBM MQ subsystems running on a single z/OS instance from a single task.
- Provides z/OS-wide coordinated commitment control (using z/OS RRS) for recoverable resources accessed through z/OS RRS-compliant recoverable managers for:
 - Applications that connect to IBM MQ using the RRS batch adapter.
 - Db2-stored procedures executing in a Db2-stored procedures address space that is managed by a workload manager (WLM) on z/OS.
- Supports the ability to switch an IBM MQ batch thread between TCBs.

IBM MQ for z/OS provides two RRS batch adapters:

CSQBRSTB

This adapter requires you to change any MQCMIT statement to SRRCMIT and any MQBACK statement to SRRBACK in your IBM MQ application. (If you code MQCMIT or MQBACK in an application linked with CSQBRSTB, you receive MQRC_ENVIRONMENT_ERROR.)

CSQBRRSI

This adapter allows your IBM MQ application to use either MQCMIT and MQBACK or SRRCMIT and SRRBACK.

Note: CSQBRSTB and CSQBRRSI are shipped with linkage attributes AMODE(31) RMODE(ANY). If your application loads either stub below the 16 MB line, first relink the stub with RMODE(24).

Migration

You can migrate existing Batch/TSO IBM MQ applications to use RRS coordination with few or no changes.

If you link-edit your IBM MQ application with the CSQBRRSI adapter, MQCMIT and MQBACK syncpoint your unit of work across IBM MQ and all other RRS-enabled resource managers. If you link-edit your IBM MQ application with the CSQBRSTB adapter, change MQCMIT to SRRCMIT and MQBACK to SRRBACK. The latter approach is preferable; it clearly indicates that the syncpoint is not restricted to IBM MQ resources only.

The IMS adapter:

If you are using the IMS adapter from an IBM MQ for z/OS system, ensure that IMS can obtain sufficient storage to accommodate messages up to 100 MB long.

Note to users

The *IMS adapter* provides access to IBM MQ for z/OS resources for:

- Online message processing programs (MPPs)
- Interactive fast path programs (IFPs)
- Batch message processing programs (BMPs)

To use these resources, the programs must be running in task (TCB) mode and problem state; they must not be in cross-memory mode or access-register mode.

The adapter provides a connection thread from an application task control block (TCB) to IBM MQ. The adapter supports a two-phase commit protocol for changes made to resources owned by IBM MQ for z/OS, with IMS acting as the syncpoint coordinator.

The adapter also provides a trigger monitor program that can start programs automatically when certain trigger conditions on a queue are met. For more information, see “Starting IBM MQ applications using triggers” on page 234.

If you are writing batch DL/I programs, follow the guidance given in this topic for z/OS batch programs.

Writing z/OS UNIX System Services applications:

The batch adapter supports queue manager connections from batch and TSO address spaces:

If we consider a Batch address space, the adapter supports connections from multiple TCBS within that address space as follows:

- Each TCB can connect to multiple queue managers using the MQCONN or MQCONNX call (but a TCB can only have one instance of a connection to a particular queue manager at any one time).
- Multiple TCBS can connect to the same queue manager (but the queue manager handle returned on any MQCONN or MQCONNX call is bound to the issuing TCB and cannot be used by any other TCB).

z/OS UNIX System Services supports two types of pthread_create call:

1. Heavyweight threads, run one for each TCB, that are ATTACHed and DETACHed at thread start and end by z/OS.

2. Medium-weight threads, run one for each TCB, but the TCB can be one of a pool of long-running TCBs. The application must perform all necessary application cleanup, because, if it is connected to a server, the default thread termination that might be provided by the server at task (TCB) termination, is **not** always driven.

Lightweight threads are not supported. (If an application creates permanent threads that dispatch their own work requests, the *application* is responsible for cleaning up any resources before starting the next work request.)

IBM MQ for z/OS supports z/OS UNIX System Services threads using the Batch Adapter as follows:

1. Heavyweight threads are fully supported as batch connections. Each thread runs in its own TCB, which is attached and detached at thread start and end. Should the thread end before issuing an MQDISC call, IBM MQ for z/OS performs its standard task cleanup, which includes committing any outstanding unit of work if the thread terminated normally, or backing it out if the thread terminated abnormally.
2. Medium-weight threads are fully supported, but if the TCB is going to be reused by another thread, the application must ensure that an MQDISC call, preceded by either MQCMIT or MQBACK, is issued before the next thread start. This implies that if the application has established a Program Interrupt Handler, and the application then abends, the Interrupt Handler must issue MQCMIT and MQDISC calls before reusing the TCB for another thread.

Note: Threading models do **not** support access to common IBM MQ resources from multiple threads.

The API-crossing exit for z/OS:

This topic contains product-sensitive programming interface information.

An exit is a point in IBM-supplied code where you can run your own code. IBM MQ for z/OS provides an *API-crossing exit* that you can use to intercept calls to the MQI, and to monitor or modify the function of the MQI calls. This section describes how to use the API-crossing exit, and describes the sample exit program that is supplied with IBM MQ for z/OS.

This section is applicable only for users of CICS TS V3.1 and earlier. Users of CICS TS V3.2 and later should refer to the section CICS Integration with IBM MQ in the CICS product documentation.

Note

The API-crossing exit is invoked only by the CICS adapter of IBM MQ for z/OS. The exit program runs in the CICS address space.

Writing your own exit program:

You can use the sample API-crossing exit program (CSQCAPX) that is supplied with IBM MQ for z/OS as a framework for your own program.

This is described in “The sample API-crossing exit program, CSQCAPX” on page 265.

When writing an exit program, to find the name of an MQI call issued by an application, examine the *ExitCommand* field of the MQXP structure. To find the number of parameters on the call, examine the *ExitParmCount* field. You can use the 16-byte *ExitUserArea* field to store the address of any dynamic storage that the application obtains. This field is retained across invocations of the exit and has the same lifetime as a CICS task.

If you are using CICS Transaction Server V3.2, you must write your exit program to be threadsafe and declare your exit program as threadsafe. If you are using earlier CICS releases, you are also recommended to write and declare your exit programs as threadsafe to be ready for migrating to CICS Transaction Server V3.2.

Your exit program can suppress execution of an MQI call by returning MQXCC_SUPPRESS_FUNCTION or MQXCC_SKIP_FUNCTION in the *ExitResponse* field. To allow the call to be executed (and the exit program to be reinvoked after the call has completed), your exit program must return MQXCC_OK.

When invoked after an MQI call, an exit program can inspect and modify the completion and reason codes set by the call.

Usage notes

Here are some general points to consider when writing your exit program:

- For performance reasons, write your program in assembler-language. If you write it in any of the other languages supported by IBM MQ for z/OS, you must provide your own data definition file.
- Link-edit your program as AMODE(31) and RMODE(ANY).
- To define the exit parameter block to your program, use the assembler-language macro, CMQXPA.
- Specify CONCURRENCY(THREADSAFE) when you define your exit program and any programs that your exit program calls.
- If you are using the CICS Transaction Server for z/OS storage protection feature, your program must run in CICS execution key. That is, you must specify EXECKEY(CICS) when defining both your exit program and any programs to which it passes control. For information about CICS exit programs and the CICS storage protection facility, see the *CICS Customization Guide*.
- Your program can use all the APIs (for example, IMS, Db2, and CICS) that a CICS task-related user exit program can use. It can also use any of the MQI calls except MQCONN, MQCONNX, and MQDISC. However, any MQI calls within the exit program do not invoke the exit program a second time.
- Your program can issue EXEC CICS SYNCPOINT or EXEC CICS SYNCPOINT ROLLBACK commands. However, these commands commit or roll back *all* the updates done by the task up to the point that the exit was used, and so their use is not recommended.
- Your program must end by issuing an EXEC CICS RETURN command. It must not transfer control with an XCTL command.
- Exits are written as extensions to the IBM MQ for z/OS code. Ensure that your exit does not disrupt any IBM MQ for z/OS programs or transactions that use the MQI. These are typically indicated with a prefix of CSQ or CK.
- If CSQCAPX is defined to CICS, the CICS system attempts to load the exit program when CICS connects to IBM MQ for z/OS. If this attempt is successful, message CSQC301I is sent to the CKQC panel or to the system console. If the load is unsuccessful (for example, if the load module does not exist in any of the libraries in the DFHRPL concatenation), message CSQC315 is sent to the CKQC panel or to the system console.
- Because the parameters in the communication area are addresses, the exit program must be defined as local to the CICS system (that is, not as a remote program).

The sample API-crossing exit program, CSQCAPX:

The sample exit program is supplied as an assembler-language program. The source file (CSQCAPX) is supplied in the library **thlqual.SCSQASMS** (where **thlqual** is the high-level qualifier used by your installation). This source file includes pseudocode that describes the program logic.

The sample program contains initialization code and a layout that you can use when writing your own exit programs.

The sample shows how to:

- Set up the exit parameter block
- Address the call and exit parameter blocks
- Determine for which MQI call the exit is being invoked
- Determine whether the exit is being invoked before or after processing of the MQI call
- Put a message on a CICS temporary storage queue
- Use the macro DFHEIENT for dynamic storage acquisition to maintain reentrancy
- Use DFHEIBLK for the CICS exec interface control block
- Trap error conditions
- Return control to the caller

Design of the sample exit program

The sample exit program writes messages to a CICS temporary storage queue (CSQ1EXIT) to show the operation of the exit.

The messages show whether the exit is being invoked before or after the MQI call. If the exit is invoked after the call, the message contains the completion code and reason code returned by the call. The sample uses named constants from the CMQXPA macro to check on the type of entry (that is, before or after the call).

The sample does not perform any monitoring function, but simply places time-stamped messages into a CICS queue indicating the type of call it is processing. This provides an indication of the performance of the MQI, as well as the correct functioning of the exit program.

Note: The sample exit program issues six EXEC CICS calls for each MQI call that is made while the program is running. If you use this exit program, IBM MQ for z/OS performance is degraded.

Preparing and using the API-crossing exit:

The sample exit is supplied in source form only.

To use the sample exit, or an exit program that you have written, create a load library, as you would for any other CICS program, as described in “Building CICS applications in z/OS” on page 425.

- For CICS Transaction Server for z/OS and CICS for MVS™/ESA, when you update the CICS system definition (CSD) data set, the definitions you need are in the member **thlqual.SCSQPROC(CSQ4B100)**.

Note: The definitions use a suffix of MQ. If this suffix is already used in your enterprise, this must be changed before the assembly stage.

If you use the default CICS program definitions supplied, the exit program CSQCAPX is installed in a *disabled* state. This is because using the exit program can produce a significant reduction in performance.

To activate the API-crossing exit temporarily:

1. Issue the command **CEMT S PROGRAM(CSQCAPX) ENABLED** from the CICS master terminal.
2. Run the CKQC transaction, and use option 3 in the Connection pull-down to alter the status of the API-crossing exit to **Enabled**.

If you want to run IBM MQ for z/OS with the API-crossing exit permanently enabled, with CICS Transaction Server for z/OS and CICS for MVS/ESA, do one of the following:

- Alter the CSQCAPX definition in member CSQ4B100, changing STATUS(DISABLED) to STATUS(ENABLED). You can update the CICS CSD definition using the CICS-supplied batch program DFHCSDUP.
- Alter the CSQCAPX definition in the CSQCAT1 group by changing the status from DISABLED to ENABLED.

In both cases, you must reinstall the group. You can do this by cold-starting your CICS system or by using the CICS CEDA transaction to reinstall the group while CICS is running.

Note: Using CEDA might cause an error if any of the entries in the group are currently in use.

End of product-sensitive programming interface information.

Application programming with shared queues:

This topic provides information on some of the factors that you need to take into account when designing new applications to use shared queues, and when migrating existing applications to the shared-queue environment.

Serializing your applications:

Certain types of applications might have to ensure that messages are retrieved from a queue in exactly the same order as they arrived on the queue.

For example, if IBM MQ is being used to shadow database updates on to a remote system, a message describing the update to a record must be processed after a message describing the insert of that record. In a local queuing environment, this is often achieved by the application that is getting the messages opening the queue with the `MQOO_INPUT_EXCLUSIVE` option, thus preventing any other getting application from processing the queue at the same time.

IBM MQ allows applications to open shared queues exclusively in the same way. However, if the application is working from a partition of a queue (for example, all database updates are on the same queue, but those for table A have a correlation identifier of A, and those for table B a correlation identifier of B), and applications want to get messages for table A updates and table B updates concurrently, the simple mechanism of opening the queue exclusively is not possible.

If this type of application is to take advantage of the high availability of shared queues, you might decide that another instance of the application that accesses the same shared queues, running on a secondary queue manager, should take over if the primary getting application or queue manager fails.

If the primary queue manager fails, two things happen:

- Shared queue peer recovery ensures that any incomplete updates from the primary application are completed or backed out.
- The secondary application takes over processing the queue.

The secondary application might start before all the incomplete units of work have been dealt with, which could lead to the secondary application retrieving the messages out of sequence. To solve this type of problem, the application can choose to be a *serialized application*.

A serialized application uses the `MQCONN` call to connect to the queue manager, specifying a connection tag when it connects that is unique to that application. Any units of work performed by the application are marked with the connection tag. IBM MQ ensures that units of work within the queue-sharing group with the same connection tag are serialized (according to the serialization options on the `MQCONN` call).

This means that, if the primary application uses the `MQCONN` call with a connection tag of Database shadow retriever, and the secondary takeover application attempts to use the `MQCONN` call with an identical connection tag, the secondary application cannot connect to the second IBM MQ until any outstanding primary units of work have been completed, in this case by peer recovery.

Consider using the serialized application technique for applications that depend on the exact sequence of messages on a queue. In particular:

- Applications that must not restart after an application or queue manager failure until all commit and backout operations for the previous execution of the application are complete.
In this case, the serialized application technique is only applicable if the application works in syncpoint.
- Applications that must not start while another instance of the same application is already running.
In this case, the serialized application technique is only required if the application cannot open the queue for exclusive input.

Note: IBM MQ only guarantees to preserve the sequence of messages when certain criteria are met. These are described in the description of `MQGET`.

Applications that are not suitable for use with shared queues:

Some features of IBM MQ are not supported when you are using shared queues, so applications that use these features are not suitable for the shared queue environment.

Consider the following points when designing your shared-queue applications:

- Queue indexing is limited for shared queues. If you want to use the message identifier or correlation identifier to select the message that you want to get from the queue, the queue should be indexed with the correct value. If you are selecting messages by message identifier alone, the queue needs an index type of MQIT_MSG_ID (although you can also use MQIT_NONE). If you are selecting messages by correlation identifier alone, the queue must have an index type of MQIT_CORREL_ID.
- You cannot use temporary dynamic queues as shared queues. However, you can use permanent dynamic queues. The models for shared dynamic queues have a DEFTYPE of SHAREDYN (shared dynamic) although they are created and destroyed in the same way as PERMDYN (permanent dynamic) queues.

Deciding whether to share non-application queues:

Use this information when considering sharing non-application queues.

There are queues other than application queues that you might want to consider sharing:

Initiation queues

If you define a shared initiation queue, you do not need to have a trigger monitor running on every queue manager in the queue-sharing group, as long as there is at least one trigger monitor running. (You can also use a shared initiation queue even if there is a trigger monitor running on each queue manager in the queue-sharing group.)

If you have a shared application queue and use the trigger type of EVERY (or a trigger type of FIRST with a small trigger interval, which behaves like a trigger type of EVERY) your initiation queue must always be a shared queue. For more information about when to use a shared initiation queue, see Table 32 on page 269.

SYSTEM.* queues

You can define the SYSTEM.ADMIN.* queues used to hold event messages as shared queues. This can be useful to check load balancing if an exception occurs. Each event message created by IBM MQ contains a correlation identifier indicating which queue manager produced it.

You must define the SYSTEM.QSG.* queues used for shared channels and intra-group queuing as shared queues.

You can also change the definitions of the SYSTEM.DEFAULT.LOCAL.QUEUE to be shared, or define your own default shared queue definition. This is described in the section **Defining system objects** in the Planning on z/OS IBM MQ for z/OS Concepts and Planning Guide.

You cannot define any other SYSTEM.* queues as shared queues.

Migrating your existing applications to use shared queues:

Reason codes, triggering, and the MQINQ API call can work differently in a shared queue environment.

Migrating your existing queues to shared queues is described in the Administering IBM MQ for z/OS *IBM MQ for z/OS System Administration Guide*.

When you migrate your existing applications, consider the following things, which might work in a different way in the shared queue environment:

Reason codes

When you migrate your existing applications to use shared queues, check for the new reason codes that can be issued.

Triggering

If you are using a shared application queue, triggering works on committed messages only (on a non-shared application queue, triggering works on all messages).

If you use triggering to start applications, you might want to use a shared initiation queue. Table 32 describes what you need to consider when deciding which type of initiation queue to use.

Table 32. When to use a shared-initiation queue

	Non-shared application queue	Shared application queue
Non-shared initiation queue	As for previous releases.	<p>If you use a trigger type of FIRST or DEPTH, you can use a non-shared initiation queue with a shared application queue. Extra trigger messages might be generated, but this setup is good for triggering long-running applications (like the CICS bridge) and provides high availability.</p> <p>For trigger type FIRST or DEPTH, a trigger message triggers an instance of the application on every queue manager that is running a trigger monitor and that does not already have the application queue open for input. One trigger message is generated for every queue manager; if there is more than one trigger monitor running against the non-shared local initiation queue, on a particular queue manager, they will compete to process the message.</p>
Shared initiation queue	Do not use a shared initiation queue with a non-shared application queue.	<p>If you have a shared application queue that has a trigger type of EVERY, use a shared initiation queue, or you might lose trigger messages in certain circumstances; for example, a queue manager failing.</p> <p>For trigger type FIRST or DEPTH, one trigger message is generated by each queue manager that has the named initiation queue open for input.</p> <p>Note: For trigger type FIRST or DEPTH, if one trigger monitor instance is busy, this leaves the potential for less busy trigger monitors to process more than one trigger message from the shared initiation queue. Hence, multiple instances of the server application may be started against a given queue manager. Note that these multiple instances are started as a result of processing multiple trigger messages. Ordinarily, for trigger type FIRST or DEPTH, if an application instance is already serving an application queue, another trigger message will not be generated by the queue manager that the application is connected to.</p>

MQINQ

When you use the MQINQ call to display information about a shared queue, the values of the number of MQOPEN calls that have the queue open for input and output relate only to the queue manager that issued the call. No information is produced about other queue managers in the queue-sharing group that have the queue open.

IMS and IMS bridge applications on IBM MQ for z/OS

z/OS

This information helps you to write IMS applications using IBM MQ.

- To use syncpoints and MQI calls in IMS applications, see “Writing IMS applications using IBM MQ” on page 58.
- To write applications that use the IBM MQ - IMS bridge, see “Writing IMS bridge applications” on page 63.

Use the following links to find out more about IMS and IMS bridge applications on IBM MQ for z/OS:

- “Writing IMS applications using IBM MQ” on page 58
- “Writing IMS bridge applications” on page 63

Related concepts:

“The Message Queue Interface overview” on page 75

Learn about the Message Queue Interface (MQI) components.

“Connecting to and disconnecting from a queue manager” on page 91

To use IBM MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

“Opening and closing objects” on page 100

This information provides an insight into opening and closing IBM MQ objects.

“Putting messages on a queue” on page 111

Use this information to learn how to put messages on a queue.

“Getting messages from a queue” on page 128

Use this information to learn about getting messages from a queue.

“Inquiring about and setting object attributes” on page 218

Attributes are the properties that define the characteristics of an IBM MQ object.

“Committing and backing out units of work” on page 222

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

“Starting IBM MQ applications using triggers” on page 234

Learn about triggers and how to start IBM MQ applications using triggers.

“Working with the MQI and clusters” on page 254

There are special options on calls and return codes that relate to clustering.

“Using and writing applications on IBM MQ for z/OS” on page 258

IBM MQ for z/OS applications can be made up from programs that run in many different environments. This means that they can take advantage of the facilities available in more than one environment.

Writing IMS applications using IBM MQ:

There are further considerations when using IBM MQ in IMS applications. These include which MQ API calls can be used and the mechanism used for syncpoint.

Use the following links to find out more about writing IMS applications on IBM MQ for z/OS:

- “Syncpoints in IMS applications” on page 59

- “MQI calls in IMS applications” on page 59

Restrictions

There are restrictions on which IBM MQ API calls can be used by an application using the IMS adapter.

The following IBM MQ API calls are not supported within an application using the IMS adapter:

- MQCB
- MQCB_FUNCTION
- MQCTL

Related concepts:

“Writing IMS bridge applications” on page 63

This topic contains information about writing applications to use the IBM MQ - IMS bridge.

Syncpoints in IMS applications:

In an IMS application, you establish a syncpoint by using IMS calls such as GU (get unique) to the IOPCB and CHKP (checkpoint).

To back out all changes since the previous checkpoint, you can use the IMS ROLB (rollback) call. For more information, see the following documentation:

- IMS Version 13 Application Programming APG SC19-3646
- IMS Version 13 Application Programming APIs APR SC19-3647

The queue manager is a participant in a two-phase commit protocol; the IMS syncpoint manager is the coordinator.

All open handles are closed by the IMS adapter at a syncpoint (except in a batch or non-message driven BMP environment). This is because a different user could initiate the next unit of work and IBM MQ security checking is performed when the MQCONN, MQCONNX, and MQOPEN calls are made, not when the MQPUT or MQGET calls are made.

However, in a Wait-for-Input (WFI) or pseudo Wait-for-Input (PWFI) environment IMS does not notify IBM MQ to close the handles until either the next message arrives or a QC status code is returned to the application. If the application is waiting in the IMS region and any of these handles belong to triggered queues, triggering will not occur because the queues are open. For this reason, applications running in a WFI or PWFI environment should explicitly MQCLOSE the queue handles before doing the GU to the IOPCB for the next message.

If an IMS application (either a BMP or an MPP) issues the MQDISC call, open queues are closed but no implicit syncpoint is taken. If the application ends normally, any open queues are closed and an implicit commit occurs. If the application ends abnormally, any open queues are closed and an implicit backout occurs.

MQI calls in IMS applications:

Use this information to learn about the use of MQI calls on Server applications and Enquiry applications.

This section covers the use of MQI calls in the following types of IMS applications:

- “Server applications” on page 59
- “Inquiry applications” on page 62

Server applications

Here is an outline of the MQI server application model:

```
Initialize/Connect
.
Open queue for input shared
.
Get message from IBM MQ queue
.
Do while Get does not fail
.
If expected message received
Process the message
Else
Process unexpected message
End if
.
Commit
.
Get next message from IBM MQ queue
.
End do
.
Close queue/Disconnect
.
END
```

Sample program CSQ4ICB3 shows the implementation, in C/370, of a BMP using this model. The program establishes communication with IMS first, and then with IBM MQ:

```
main()
----
Call InitIMS
If IMS initialization successful
Call InitMQM
If IBM MQ initialization successful
Call ProcessRequests
Call EndMQM
End-if
End-if

Return
```

The IMS initialization determines whether the program has been called as a message-driven or a batch-oriented BMP and controls IBM MQ queue manager connection and queue handles accordingly:

```
InitIMS
-----
Get the IO, Alternate and Database PCBs
Set MessageOriented to true

Call ctdli to handle status codes rather than abend
If call is successful (status code is zero)
While status code is zero
Call ctdli to get next message from IMS message queue
If message received
```



```

Do nothing
Else if no IOPBC
Set MessageOriented to false
Initialize error message
Build 'Started as batch oriented BMP' message
Call ReportCallError to output the message
End-if
Else if response is not 'no message available'
Initialize error message
Build 'GU failed' message
Call ReportCallError to output the message
Set return code to error
End-if
End-if
End-while
Else
Initialize error message
Build 'INIT failed' message
Call ReportCallError to output the message
Set return code to error
End-if

```

Return to calling function

The IBM MQ initialization connects to the queue manager and opens the queues. In a message-driven BMP this is called after each IMS syncpoint is taken; in a batch-oriented BMP, this is called only during program startup:

```

InitMQM
-----
Connect to the queue manager
If connect is successful
Initialize variables for the open call
Open the request queue
If open is not successful
Initialize error message
Build 'open failed' message
Call ReportCallError to output the message
Set return code to error
End-if
Else
Initialize error message
Build 'connect failed' message
Call ReportCallError to output the message
Set return code to error
End-if

```

Return to calling function

The implementation of the server model in an MPP is influenced by the fact that the MPP processes a single unit of work per invocation. This is because, when a syncpoint (GU) is taken, the connection and queue handles are closed and the next IMS message is delivered. This limitation can be partially overcome by one of the following:

- **Processing many messages within a single unit-of-work**

This involves:

- Reading a message
- Processing the required updates
- Putting the reply

in a loop until all messages have been processed or until a set maximum number of messages has been processed, at which time a syncpoint is taken.

Only certain types of application (for example, a simple database update or inquiry) can be approached in this way. Although the MQI reply messages can be put with the authority of the originator of the MQI message being handled, the security implications of any IMS resource updates need to be addressed carefully.

- **Processing one message per invocation of the MPP and ensuring multiple scheduling of the MPP to process all available messages.**

Use the IBM MQ IMS trigger monitor program (CSQQTRMN) to schedule the MPP transaction when there are messages on the IBM MQ queue and no applications serving it.

If trigger monitor starts the MPP, the queue manager name and queue name are passed to the program, as shown in the following COBOL code extract:

```
* Data definition extract
01 WS-INPUT-MSG.
05 IN-LL1          PIC S9(3) COMP.
05 IN-ZZ1          PIC S9(3) COMP.
05 WS-STRINGPARM  PIC X(1000).
01 TRIGGER-MESSAGE.
COPY CMQTMCL.
*
* Code extract
GU-IOPCB SECTION.
MOVE SPACES TO WS-STRINGPARM.
CALL 'CBLTDLI' USING GU,
IOPCB,
WS-INPUT-MSG.
IF IOPCB-STATUS = SPACES
MOVE WS-STRINGPARM TO MQTMC.
* ELSE handle error
*
* Now use the queue manager and queue names passed
DISPLAY 'MQTMC-QMGRNAME   ='
MQTMC-QMGRNAME OF MQTMC '='.
DISPLAY 'MQTMC-QNAME     ='
MQTMC-QNAME   OF MQTMC '='.
```

The server model, which is expected to be a long running task, is better supported in a batch processing region, although the BMP cannot be triggered using CSQQTRMN.

Inquiry applications

A typical IBM MQ application initiating an inquiry or update works as follows:

- Gather data from the user
- Put one or more IBM MQ messages
- Get the reply messages (you might have to wait for them)
- Provide a response to the user

Because messages put on to IBM MQ queues do not become available to other IBM MQ applications until they are committed, they must either be put out of syncpoint, or the IMS application must be split into two transactions.

If the inquiry involves putting a single message, you can use the *no syncpoint* option; however, if the inquiry is more complex, or resource updates are involved, you might get consistency problems if failure occurs and you do not use syncpointing.

To overcome this, you can split IMS MPP transactions using MQI calls using a program-to-program message switch; see *IMS/ESA Application Programming: Data Communication* for information about this. This allows an inquiry program to be implemented in an MPP:

```
Initialize first program/Connect
.
Open queue for output
```

```

.
Put inquiry to IBM MQ queue
.
Switch to second IBM MQ program, passing necessary data in save
pack area (this commits the put)
.
END
.
.
Initialize second program/Connect
.
Open queue for input shared
.
Get results of inquiry from IBM MQ queue
.
Return results to originator
.
END

```

Writing IMS bridge applications:

This topic contains information about writing applications to use the IBM MQ - IMS bridge.

For information about the IBM MQ - IMS bridge, see [The IMS bridge](#).

Use the following links to find out more about writing IMS bridge applications on IBM MQ for z/OS:

- “How the IMS bridge deals with messages” on page 63
- “Writing IMS transaction programs through IBM MQ” on page 283

Related concepts:

“Writing IMS applications using IBM MQ” on page 58

There are further considerations when using IBM MQ in IMS applications These include which MQ API calls can be used and the mechanism used for syncpoint.

How the IMS bridge deals with messages:

When you use the IBM MQ - IMS bridge to send messages to an IMS application, you need to construct your messages in a special format.

You must also put your messages on IBM MQ queues that have been defined with a storage class that specifies the XCF group and member name of the target IMS system. These are known as MQ-IMS bridge queues, or simply **bridge** queues.

The IBM MQ-IMS bridge requires exclusive input access (MQOO_INPUT_EXCLUSIVE) to the bridge queue if it is defined with QSGDISP(QMGR), or if it is defined with QSGDISP(SHARED) together with the NOSHARE option.

A user does not need to sign on to IMS before sending messages to an IMS application. The user ID in the *UserIdentifier* field of the MQMD structure is used for security checking. The level of checking is determined when IBM MQ connects to IMS, and is described in [Application access control for the IMS bridge](#). This enables a pseudo signon to be implemented.

The IBM MQ - IMS bridge accepts the following types of message:

- Messages containing IMS transaction data and an MQIIH structure (described in [MQIIH](#)):
MQIIH LLZZ<trancode><data>[LLZZ<data>] [LLZZ<data>]

Note:

1. The square brackets, [], represent optional multi-segments.

- 2. Set the *Format* field of the MQMD structure to MQFMT_IMS to use the MQIIH structure.
- Messages containing IMS transaction data but no MQIIH structure:

```
LLZZ<trancode><data> \
[LLZZ<data>][LLZZ<data>]
```

IBM MQ validates the message data to ensure that the sum of the LL bytes plus the length of the MQIIH (if it is present) is equal to the message length.

When the IBM MQ - IMS bridge gets messages from the bridge queues, it processes them as follows:

- If the message contains an MQIIH structure, the bridge verifies the MQIIH (see MQIIH), builds the OTMA headers, and sends the message to IMS. The transaction code is specified in the input message. If this is an LTERM, IMS replies with a DFS1288E message. If the transaction code represents a command, IMS executes the command; otherwise the message is queued in IMS for the transaction.
- If the message contains IMS transaction data, but no MQIIH structure, the IMS bridge makes the following assumptions:
 - The transaction code is in bytes 5 through 12 of the user data
 - The transaction is in nonconversational mode
 - The transaction is in commit mode 0 (commit-then-send)
 - The *Format* in the MQMD is used as the *MFSMapName* (on input)
 - The security mode is MQISS_CHECK

The reply message is also built without an MQIIH structure, taking the *Format* for the MQMD from the *MFSMapName* of the IMS output.

The IBM MQ - IMS bridge uses one or two Tpipes for each IBM MQ queue:

- A synchronized Tpipe is used for all messages using Commit mode 0 (COMMIT_THEN_SEND) (these show with SYN in the status field of the IMS /DIS TMEMBER client TPIPE xxxx command)
- A non-synchronized Tpipe is used for all messages using Commit mode 1 (SEND_THEN_COMMIT)

The Tpipes are created by IBM MQ when they are first used. A non-synchronized Tpipe exists until IMS is restarted. Synchronized Tpipes exist until IMS is cold started. You cannot delete these Tpipes yourself.

See the following topics for more information about how the IBM MQ - IMS bridge deals with messages:

- “Mapping IBM MQ messages to IMS transaction types” on page 64
- “If the message cannot be put to the IMS queue” on page 65
- “IMS bridge feedback codes” on page 65
- “The MQMD fields in messages from the IMS bridge” on page 66
- “The MQIIH fields in messages from the IMS bridge” on page 67
- “Reply messages from IMS” on page 68
- “Using alternate response PCBs in IMS transactions” on page 68
- “Sending unsolicited messages from IMS” on page 68
- “Message segmentation” on page 69
- “Data conversion” on page 69

Related concepts:

“Writing IMS transaction programs through IBM MQ” on page 283

The coding required to handle IMS transactions through IBM MQ depends on the message format required by the IMS transaction and the range of responses it can return. However, there are several points to consider when your application handles IMS screen formatting information.

Mapping IBM MQ messages to IMS transaction types:

A table describing the mapping of IBM MQ messages to IMS transaction types.

Table 33. Mapping IBM MQ messages to IMS transaction types

IBM MQ message type	Commit-then-send (mode 0) - uses synchronized IMS Tpipes	Send-then-commit (mode 1) - uses non-synchronized IMS Tpipes
Persistent IBM MQ messages	<ul style="list-style-type: none">Recoverable full function transactionsUnrecoverable transactions are rejected by IMS	<ul style="list-style-type: none">Fastpath transactionsConversational transactionsFull function transactions
Nonpersistent IBM MQ messages	<ul style="list-style-type: none">Unrecoverable full function transactionsRecoverable transactions are permitted with IMS V8 and APAR PQ61404 and all later versions of IMS	<ul style="list-style-type: none">Fastpath transactionsConversational transactionsFull function transactions

Note: IMS commands cannot use persistent IBM MQ messages with commit mode 0. See the *IMS/ESA Open Transaction Manager Access User's Guide* for more information.

If the message cannot be put to the IMS queue:

Learn about actions to take if the message cannot be put to the IMS queue.

If the message cannot be put to the IMS queue, the following action is taken by IBM MQ:

- If a message cannot be put to IMS because the message is invalid, the message is put to the dead-letter queue, and a message is sent to the system console.
- If the message is valid, but is rejected by IMS, IBM MQ sends an error message to the system console, the message includes the IMS sense code, and the IBM MQ message is put to the dead-letter queue. If the IMS sense code is 001A, IMS sends an IBM MQ message containing the reason for the failure to the reply-to queue.

Note: In the circumstances listed previously, if IBM MQ cannot put the message to the dead-letter queue for any reason, the message is returned to the originating IBM MQ queue. An error message is sent to the system console, and no further messages are sent from that queue.

To resend the messages, do *one* of the following:

- Stop and restart the Tpipes in IMS corresponding to the queue
 - Alter the queue to GET(DISABLED), and again to GET(ENABLED)
 - Stop and restart IMS or the OTMA
 - Stop and restart your IBM MQ subsystem
- If the message is rejected by IMS for anything other than a message error, the IBM MQ message is returned to the originating queue, IBM MQ stops processing the queue, and an error message is sent to the system console.

If an exception report message is required, the bridge puts it to the reply-to queue with the authority of the originator. If the message cannot be put to the queue, the report message is put to the dead-letter queue with the authority of the bridge. If it cannot be put to the DLQ, it is discarded.

IMS bridge feedback codes:

IMS sense codes are typically output in hexadecimal format in IBM MQ console messages such as CSQ2001I (for example, sense code 0x001F). IBM MQ feedback codes as seen in the dead-letter header of messages put to the dead-letter queue are decimal numbers.

The IMS bridge feedback codes are in the range 301 through 399, or 600 through 855 for NACK sense code 0x001A. They are mapped from the IMS-OTMA sense codes as follows:

1. The IMS-OTMA sense code is converted from a hexadecimal number to a decimal number.
2. 300 is added to the number resulting from the calculation in 1, giving the IBM MQ *Feedback* code.
3. The IMS-OTMA sense code 0x001A, decimal 26 is a special case. A *Feedback* code in the range 600-855 is generated.
 - a. The IMS-OTMA reason code is converted from a hexadecimal number to a decimal number.
 - b. 600 is added to the number resulting from the calculation in a, giving the IBM MQ *Feedback* code.

For information about IMS-OTMA sense codes, see the IMS Messages and Codes.

The MQMD fields in messages from the IMS bridge:

Learn about the MQMD fields in messages from the IMS bridge.

The MQMD of the originating message is carried by IMS in the User Data section of the OTMA headers. If the message originates in IMS, this is built by the IMS Destination Resolution Exit. The MQMD of a message received from IMS is built as follows:

StrucID

"MD "

Version

MQMD_VERSION_1

Report

MQRO_NONE

MsgType

MQMT_REPLY

Expiry If MQIIH_PASS_EXPIRATION is set in the Flags field of the MQIIH, this field contains the remaining expiry time, else it is set to MQEI_UNLIMITED

Feedback

MQFB_NONE

Encoding

MQENC.Native (the encoding of the z/OS system)

CodedCharSetId

MQCCSI_Q_MGR (the CodedCharSetID of the z/OS system)

Format

MQFMT_IMS if the MQMD.Format of the input message is MQFMT_IMS, otherwise IOPCB.MODNAME

Priority

MQMD.Priority of the input message

Persistence

Depends on commit mode: MQMD.Persistence of the input message if CM-1; persistence matches recoverability of the IMS message if CM-0

MsgId MQMD.MsgId if MQRO_PASS_MSG_ID, otherwise New MsgId (the default)

CorrelId

MQMD.CorrelId from the input message if MQRO_PASS_CORREL_ID, otherwise MQMD.MsgId from the input message (the default)

BackoutCount

0

ReplyToQ

Blanks

ReplyToQMgr

Blanks (set to local qmgr name by the queue manager during the MQPUT)

UserIdentifier

MQMD.UserIdentifier of the input message

AccountingToken

MQMD.AccountingToken of the input message

ApplIdentityData

MQMD.ApplIdentityData of the input message

PutApplType

MQAT_XCF if no error, otherwise MQAT_BRIDGE

PutApplName

<XCFgroupName><XCFmemberName> if no error, otherwise QMGR name

PutDate

Date when message was put

PutTime

Time when message was put

ApplOriginData

Blanks

The MQIIH fields in messages from the IMS bridge:

Learn about the MQIIH fields in messages from the IMS bridge.

The MQIIH of a message received from IMS is built as follows:

StrucId

"IIH "

Version

1

StrucLength

84

Encoding

MQENC_NATIVE

CodedCharSetId

MQCCSI_Q_MGR

Format

MQIIH.ReplyToFormat of the input message if MQIIH.ReplyToFormat is not blank, otherwise IOPCB.MODNAME

Flags 0

LTermOverride

LTERM name (Tpipe) from OTMA header

MFSMapName

Map name from OTMA header

ReplyToFormat

Blanks

Authenticator

MQIIH.Authenticator of the input message if the reply message is being put to an MQ-IMS bridge queue, otherwise blanks.

TranInstanceId

Conversation ID / Server Token from OTMA header if in conversation, otherwise nulls

TranState

"C" if in conversation, otherwise blank

CommitMode

Commit mode from OTMA header ("0" or "1")

SecurityScope

Blank

Reserved

Blank

Reply messages from IMS:

When an IMS transaction ISRTs to its IOPCB, the message is routed back to the originating LTERM or TPIPE.

These are seen in IBM MQ as reply messages. Reply messages from IMS are put onto the reply-to queue specified in the original message. If the message cannot be put onto the reply-to queue, it is put onto the dead-letter queue using the authority of the bridge. If the message cannot be put onto the dead-letter queue, a negative acknowledgment is sent to IMS to say that the message cannot be received.

Responsibility for the message is then returned to IMS. If you are using commit mode 0, messages from that Tpipe are not sent to the bridge, and remain on the IMS queue; that is, no further messages are sent until restart. If you are using commit mode 1, other work can continue.

If the reply has an MQIIH structure, its format type is MQFMT_IMS; if not, its format type is specified by the IMS MOD name used when inserting the message.

Using alternate response PCBs in IMS transactions:

When an IMS transaction uses alternate response PCBs (ISRTs to the ALTPCB, or issues a CHNG call to a modifiable PCB), the pre-routing exit (DFSYPRX0) is invoked to determine if the message should be rerouted.

If the message is to be rerouted, the destination resolution exit (DFSYDRU0) is invoked to confirm the destination and prepare the header information. See *Using OTMA exits in IMS* and *The pre-routing exit DFSYPRX0* for information about these exit programs.

Unless action is taken in the exits, all output from IMS transactions initiated from an IBM MQ queue manager, whether to the IOPCB or to an ALTPCB, will be returned to the same queue manager.

Sending unsolicited messages from IMS:

To send messages from IMS to an IBM MQ queue, you need to invoke an IMS transaction that ISRTs to an ALTPCB.

You need to write pre-routing and destination resolution exits to route unsolicited messages from IMS and build the OTMA user data, so that the MQMD of the message can be built correctly. See *The pre-routing exit DFSYPRX0* and *The destination resolution user exit* for information about these exit programs.

Note: The IBM MQ - IMS bridge does not know whether a message that it receives is a reply or an unsolicited message. It handles the message the same way in each case, building the MQMD and MQIIH of the reply based on the OTMA UserData that arrived with the message.

Unsolicited messages can create new Tpipes. For example, if an existing IMS transaction switched to a new LTERM (for example PRINT01), but the implementation requires that the output be delivered through OTMA, a new Tpipe (called PRINT01 in this example) is created. By default, this is a non-synchronized Tpipe. If the implementation requires the message to be recoverable, set the destination resolution exit output flag. See the *IMS Customization Guide* for more information.

Message segmentation:

You can define IMS transactions as expecting single- or multi-segment input.

The originating IBM MQ application must construct the user input following the MQIIH structure as one or more LLZZ-data segments. All segments of an IMS message must be contained in a single IBM MQ message sent with a single MQPUT.

The maximum length of an LLZZ-data segment is defined by IMS/OTMA (32767 bytes). The total IBM MQ message length is the sum of the LL bytes, plus the length of the MQIIH structure.

All the segments of the reply are contained in a single IBM MQ message.

There is a further restriction on the 32 KB limitation on messages with format MQFMT_IMS_VAR_STRING. When the data in an ASCII-mixed CCSID message is converted to an EBCDIC-mixed CCSID message, a shift-in byte or a shift-out byte is added every time that there is a transition between SBCS and DBCS characters. The 32 KB restriction applies to the maximum size of the message. That is, because the LL field in the message cannot exceed 32 KB, the message must not exceed 32 KB including all shift-in and shift-out characters. The application building the message must allow for this.

Data conversion:

The data conversion is performed by either the distributed queuing facility (which may call any necessary exits) or by the intra group queuing agent (which does not support the use of exits) when it puts a message to a destination queue that has XCF information defined for its storage class.

Any exits needed must be available to the distributed queuing facility in the data set referenced by the CSQXLIB DD statement. This means that you can send messages to an IMS application using the IBM MQ - IMS bridge from any IBM MQ platform.

If there are conversion errors, the message is put to the queue unconverted; this results eventually in it being treated as an error by the IBM MQ - IMS bridge, because the bridge cannot recognize the header format. If a conversion error occurs, an error message is sent to the z/OS console.

See "Writing data-conversion exits" on page 364 for detailed information about data conversion in general.

Sending messages to the IBM MQ - IMS bridge

To ensure that conversion is performed correctly, you must tell the queue manager what the format of the message is.

If the message has an MQIIH structure, the *Format* in the MQMD must be set to the built-in format MQFMT_IMS, and the *Format* in the MQIIH must be set to the name of the format that describes your message data. If there is no MQIIH, set the *Format* in the MQMD to your format name.

If your data (other than the LLZZs) is all character data (MQCHAR), use as your format name (in the MQIIH or MQMD, as appropriate) the built-in format MQFMT_IMS_VAR_STRING. Otherwise, use your own format name, in which case you must also provide a data-conversion exit for your format. The exit must handle the conversion of the LLZZs in your message, in addition to the data itself (but it does not have to handle any MQIIH at the start of the message).

If your application uses *MFSMapName*, you can use messages with the MQFMT_IMS instead, and define the map name passed to the IMS transaction in the MFSMapName field of the MQIIH.

Receiving messages from the IBM MQ - IMS bridge

If an MQIIH structure is present on the original message that you are sending to IMS, one is also present on the reply message.

To ensure that your reply is converted correctly:

- If you have an MQIIH structure on your original message, specify the format that you want for your reply message in the MQIIH *ReplytoFormat* field of the original message. This value is placed in the MQIIH *Format* field of the reply message. This is particularly useful if all your output data is of the form LLZZ<character data>.
- If you do not have an MQIIH structure on your original message, specify the format that you want for the reply message as the MFS MOD name in the IMS application's ISRT to the IOPCB.

Writing IMS transaction programs through IBM MQ:

The coding required to handle IMS transactions through IBM MQ depends on the message format required by the IMS transaction and the range of responses it can return. However, there are several points to consider when your application handles IMS screen formatting information.

When an IMS transaction is started from a 3270 screen, the message passes through IMS Message Format Services. This can remove all terminal dependency from the data stream seen by the transaction. When a transaction is started through OTMA, MFS is not involved. If application logic is implemented in MFS, this must be re-created in the new application.

In some IMS transactions, the end-user application can modify certain 3270 screen behavior, for example, highlighting a field that has had invalid data entered. This type of information is communicated by adding a two-byte attribute field to the IMS message for each screen field that needs to be modified by the program.

Thus, if you are coding an application to mimic a 3270, you need to take account of these fields when building or receiving messages.

You might need to code information in your program to process:

- Which key is pressed (for example, Enter and PF1)
- Where the cursor is when the message is passed to your application
- Whether the attribute fields have been set by the IMS application
 - High, normal, or zero intensity
 - Color
 - Whether IMS is expecting the field back the next time that Enter is pressed
- Whether the IMS application has used null characters (X'3F') in any fields.

If your IMS message contains only character data (apart from the LLZZ-data segment), and you are using an MQIIH structure, set the MQMD format to MQFMT_IMS and the MQIIH format to MQFMT_IMS_VAR_STRING.

If your IMS message contains only character data (apart from the LLZZ-data segment), and you are **not** using an MQIIH structure, set the MQMD format to MQFMT_IMS_VAR_STRING and ensure that your IMS application specifies MODname MQFMT_IMS_VAR_STRING when replying. If a problem occurs (for example, user not authorized to use the transaction) and IMS sends an error message, this has an MODname of the form DFSMOx, where x is a number in the range 1 through 5. This is put in the MQMD.Format.

If your IMS message contains binary, packed, or floating point data (apart from the LLZZ-data segment), code your own data-conversion routines. Refer to *IMS/ESA Application Programming: Transaction Manager* for information about IMS screen formatting.

Consider the following topics when writing code to handle IMS transactions through IBM MQ.

- “Writing IBM MQ applications to invoke IMS conversational transactions”
- “Writing programs containing IMS commands” on page 284
- “Triggering” on page 284

Writing IBM MQ applications to invoke IMS conversational transactions

Use this information as a guide for considerations when writing IBM MQ application to invoke IMS conversational transactions.

When you write an application that invokes an IMS conversation, consider the following:

- Include an MQIHH structure with your application message.
- Set the *CommitMode* in MQIHH to MQICM_SEND_THEN_COMMIT.
- To invoke a new conversation, set *TranState* in MQIHH to MQITS_NOT_IN_CONVERSATION.
- To invoke second and subsequent steps of a conversation, set *TranState* to MQITS_IN_CONVERSATION, and set *TranInstanceId* to the value of that field returned in the previous step of the conversation.
- There is no easy way in IMS to find the value of a *TranInstanceId*, should you lose the original message sent from IMS.
- The application must check the *TranState* of messages from IMS to check whether the IMS transaction has terminated the conversation.
- You can use /EXIT to end a conversation. You must also quote the *TranInstanceId*, set *TranState* to MQITS_IN_CONVERSATION, and use the IBM MQ queue on which the conversation is being carried out.
- You cannot use /HOLD or /REL to hold or release a conversation.
- Conversations invoked through the IBM MQ - IMS bridge are terminated if IMS is restarted.

Writing programs containing IMS commands

An application program can build an IBM MQ message of the form LLZZ*command*, instead of a transaction, where *command* is of the form /DIS TRAN PART or /DIS POOL ALL.

Most IMS commands can be issued in this way; see *IMS V11 Communications and Connections* for details. The command output is received in the IBM MQ reply message in the text form as would be sent to a 3270 terminal for display.

OTMA has implemented a special form of the IMS display transaction command, which returns an architected form of the output. The exact format is defined in *IMS V11 Communications and Connections*. To invoke this form from an IBM MQ message, build the message data as before, for example /DIS TRAN PART, and set the *TranState* field in the MQIHH to MQITS_ARCHITECTED. IMS processes the command, and returns the reply in the architected form. An architected response contains all the information that could be found in the text form of the output, and one additional piece of information: whether the transaction is defined as recoverable or non-recoverable.

Triggering

The IBM MQ - IMS bridge does not support trigger messages.

If you define an initiation queue that uses a storage class with XCF parameters, messages put to that queue are rejected when they get to the bridge.

Writing client procedural applications

What you need to know to write client applications on IBM MQ using a procedural language.

Applications can be built and run in the IBM MQ client environment. The application must be built and linked to the IBM MQ MQI client used. The way in which applications are built and linked varies according to the platform and programming language used. For information on how to build client applications, see “Building applications for IBM MQ MQI clients” on page 291.

You can run an IBM MQ application both in a full IBM MQ environment and in an IBM MQ MQI client environment without changing your code, provided that certain conditions are met. For more information on running your applications in the IBM MQ client environment, see “Running applications in the IBM MQ MQI client environment” on page 294.

If you use the message queue interface (MQI) to write applications to run in an IBM MQ MQI client environment there are some additional controls to impose during an MQI call to ensure that the IBM MQ application processing is not disrupted. For more information about these controls, see “Using the MQI in a client application” on page 286.

See the following topics for information preparing and running other application types as client applications:

- “Preparing and running CICS and Tuxedo applications” on page 308
- “Preparing and running Microsoft Transaction Server applications” on page 41
- “Preparing and running IBM MQ JMS applications” on page 311

Related concepts:

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

“Developing applications for IBM MQ” on page 1

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM MQ support applications written in procedural languages, and object oriented languages and frameworks.

“Design considerations for IBM MQ applications” on page 43

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

“Writing a procedural application for queuing” on page 74

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

“Writing publish/subscribe applications” on page 172

Start writing publish/subscribe IBM MQ applications.

“Building a procedural application” on page 382

You can write an IBM MQ application in one of several procedural languages, and run the application on several different platforms.

“Handling procedural program errors” on page 436

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Related tasks:

“Using the IBM MQ sample procedural programs” on page 457

These sample programs are written in procedural languages, and demonstrate typical uses of the Message Queue Interface (MQI). IBM MQ programs on different platforms.

“Developing web services with IBM MQ” on page 1217

You can develop IBM MQ applications for web services using the IBM MQ transport for SOAP.

Using the MQI in a client application

This collection of topics considers the differences between writing your IBM MQ application to run in a message queue interface (MQI) client environment and to run in the full IBM MQ queue manager environment.


When you design an application, consider what controls you need to impose during an MQI call to ensure that the IBM MQ application processing is not disrupted.

Before you can run applications that use the MQI you must create certain IBM MQ objects. For more information, see [Application programs using the MQI](#).

Limiting the size of a message in a client application:

A queue manager has a maximum message length, but the maximum size of message you can transmit from a client application is limited by the channel definition.

The maximum message length (MaxMsgLength) attribute of a queue manager is the maximum length of a message that can be handled by that queue manager.

 On Multiplatforms, you can increase the maximum message length attribute of a queue manager. For more information, see [ALTER QMGR](#).

You can find out the value of MaxMsgLength for a queue manager by using the MQINQ call.

If the MaxMsgLength attribute is changed, no check is made that there are not already queues, and even messages, with a length greater than the new value. After you change this attribute, restart applications and channels in order to ensure that the change has taken effect. It is then not possible for any new messages to be generated that exceed the MaxMsgLength of either the queue manager or the queue (unless queue manager segmentation is allowed).

The maximum message length in a channel definition limits the size of a message that you can transmit along a client connection. If an IBM MQ application tries to use the MQPUT call or the MQGET call with a message larger than this, an error code is returned to the application. The maximum message size parameter of the channel definition does not affect the maximum message size which can be consumed using MQCB over a client connection.

Related concepts:

“Using MQCONN” on page 291

You can use the MQCONN call to specify a channel definition (MQCD) structure in the MQCNO structure.

Related information:

Maximum message length (MAXMSGL)

ALTER CHANNEL

2010 (07DA) (RC2010): MQRC_DATA_LENGTH_ERROR

Choosing client or server CCSID:

Use the local coded character set identifier (CCSID) for the client. The queue manager performs necessary conversion. Use the MQCCSID environment variable to override the CCSID. If your application performs multiple PUTs, the CCSID and encoding fields of the MQMD can be overwritten after completion of the first PUT.

The data passed across the message queue interface (MQI) from the application to the client stub must be in the local CCSID, encoded for the IBM MQ MQI client. If the connected queue manager requires the data to be converted, then conversion is done by the client support code on the queue manager.

In IBM WebSphere MQ Version 7.0 and later versions, the Java client can do the conversion if the queue manager is unable to do so. See “IBM MQ classes for Java client connections” on page 1044.

The client code assumes that the character data crossing the MQI in the client is in the CCSID configured for that workstation. If this CCSID is an unsupported CCSID or is not the required CCSID, it can be overridden with the MQCCSID environment variable by using one of these commands:

- **Windows**
SET MQCCSID=850
- **UNIX**
export MQCCSID=850
- **IBM i**
ADDENVVAR ENVVAR(MQCCSID) VALUE(37)

If this parameter is set in the profile, all MQI data is assumed to be in code page 850.

Note: The assumption about code page 850 does not apply to application data in the message.

If your application is performing multiple PUTs that include IBM MQ headers after the message descriptor (MQMD), be aware that the CCSID and encoding fields of the MQMD are overwritten after completion of the first PUT.

After the first PUT, these fields contain the value used by the connected queue manager to convert the IBM MQ headers. Ensure that your application resets the values to the values it requires.

Using MQINQ in a client application:

Some values queried using MQINQ are modified by the client code.

CCSID

is set to the client CCSID, not that of the queue manager.

MaxMsgLength

is reduced if it is restricted by the channel definition. This will be the lower of:

- The value defined in the queue definition, or
- The value defined in the channel definition

For more information, see the MQINQ.

Using sync point coordination in a client application:

An application running on the base client can issue MQCMIT and MQBACK, but the scope of the sync point control is limited to the MQI resources. You can use an external transaction manager with an extended transactional client.

Within IBM MQ, one of the roles of the queue manager is sync point control within an application. If an application runs on an IBM MQ base client, it can issue MQCMIT and MQBACK, but the scope of the sync point control is limited to the MQI resources. The IBM MQ verb MQBEGIN is not valid in a base client environment.

Applications running in the full queue manager environment on the server can coordinate multiple resources (for example databases) via a transaction monitor. On the server you can use the Transaction Monitor supplied with IBM MQ products, or another transaction monitor such as CICS. You cannot use a transaction monitor with a base client application.

You can use an external transaction manager with an IBM MQ extended transactional client. See [What is an extended transactional client?](#) for details.

Using read ahead in a client application:

You can use read ahead on a client to allow non persistent messages to be sent to a client without the client application having to request the messages.

When a client requires a message from a server, it sends a request to the server. It sends a separate request for each of the messages it consumes. To improve the performance of a client consuming non persistent messages by avoiding having to send these request messages, a client can be configured to use read ahead. Read ahead allows messages to be sent to a client without an application having to request them.

Using read ahead can improve performance when consuming non persistent messages from a client application. This performance improvement is available to both MQI and JMS applications. Client applications using MQGET or asynchronous consumption benefit from the performance improvements when consuming non persistent messages.

When you call MQOPEN with MQOO_READ_AHEAD, the IBM MQ client only enables read-ahead if certain conditions are met. These conditions include:

- Both the client and remote queue manager must be at WebSphere MQ Version 7 or later.
- The client application must be compiled and linked against the threaded IBM MQ MQI client libraries.
- The client channel must be using TCP/IP protocol
- The channel must have a non-zero SharingConversations (SHARECNV) setting in both the client and server channel definitions.

When read ahead is enabled, messages are sent to a memory buffer on the client called the read ahead buffer. The client has a read ahead buffer for each queue it has open with read ahead enabled. The messages in the read ahead buffer are not persisted. The client periodically updates the server with information about the amount of data it has consumed.

Not all client application designs are suited to using read ahead because not all options are supported for use. Some options are required to be consistent between MQGET calls when read ahead is enabled. If a client alters its selection criteria between MQGET calls, messages being stored in the read ahead buffer remain stranded in the client read ahead buffer. For more information, see “Improving performance of non-persistent messages” on page 147

Read ahead configuration is controlled by three attributes, MaximumSize, PurgeTime, and UpdatePercentage, which are specified in the MessageBuffer stanza of the IBM MQ client configuration file.

Using asynchronous put in a client application:

Using asynchronous put, an application can put a message to a queue without waiting for a response from the queue manager. You can use this to improve messaging performance in some situations.

Normally, when an application puts a message or messages on a queue, using MQPUT or MQPUT1, the application has to wait for the queue manager to confirm that it has processed the MQI request. You can improve messaging performance, particularly for applications that use client bindings, and applications that put large numbers of small messages to a queue, by choosing instead to put messages asynchronously. When an application puts a message asynchronously, the queue manager does not return the success or failure of each call, but you can instead check for errors periodically.

To put a message on a queue asynchronously, use the MQPMO_ASYNC_RESPONSE option in the *Options* field of the MQPMO structure.

If a message is not eligible for asynchronous put, it is put to a queue synchronously.

When requesting asynchronous put response for MQPUT or MQPUT1, a CompCode and Reason of MQCC_OK and MQRC_NONE does not necessarily mean that the message was successfully put to a queue. Although the success or failure of each individual MQPUT or MQPUT1 call might not be returned immediately, the first error that occurred under an asynchronous call can be determined later through a call to MQSTAT.

For more details on MQPMO_ASYNC_RESPONSE, see MQPMO options.

The Asynchronous Put sample program demonstrates some of the features available. For details of the features and design of the program, and how to run it, see “The Asynchronous Put sample program” on page 477.

Using sharing conversations in a client application:

In an environment where sharing conversations is permitted, conversations can share an MQI channel instance.

Sharing conversations is controlled by two fields, both called SharingConversations, one of which is part of the channel definition (MQCD) structure and one of which is part of the channel exit parameter (MQCXP) structure. The SharingConversations field in the MQCD is an integer value, determining the maximum number of conversations that can share a channel instance associated with the channel. The SharingConversations field in the MQCXP is a boolean value, indicating whether the channel instance is currently shared.

In an environment where sharing conversations is not permitted, new client connections specifying identical MQCDs will not share a channel instance.

A new client application connection will share the channel instance when the following conditions are true:

- Both the client-connection and server-connection ends of the channel instance are configured for sharing conversations, and these values are not overridden by channel exits.
- The client connection MQCD value (supplied on the client MQCONN call or from the client channel definition table (CCDT)) exactly matches the client connection MQCD value supplied on the client MQCONN call or from the CCDT when the existing channel instance was first established. Note that the original MQCD might have been subsequently altered by exits or by channel negotiation, but that the match is made against the value which was supplied to the client system before these changes were made.
- The sharing conversations limit on the server side is not exceeded.

If a new client application connection matches the criteria to run sharing a channel instance with other conversations, this decision is made before any exits are called on that conversation. Exits on such a conversation cannot alter the fact that it is sharing the channel instance with other conversations. If there are no existing channel instances matching the new channel definition, a new channel instance is connected.

Channel negotiation only occurs for the first conversation on a channel instance; the negotiated values for the channel instance are fixed at that stage and cannot be altered when subsequent conversations start. TLS authentication also only occurs for the first conversation.

If the MQCD SharingConversations value is altered during the initialization of any security, send or receive exits for the first conversation on the socket at either the client-connection or the server-connection end of the channel instance, the new value it has after all these exits are initialized is used to determine the sharing conversations value for the channel instance (the lowest value takes precedence).

If the negotiated value for sharing conversations is zero, the channel instance is never shared. Further exit programs that set this field to zero similarly run on their own channel instance.

If the negotiated value for sharing conversations is greater than zero then MQCXP SharingConversations is set to TRUE for subsequent calls to exits, indicating that other exit programs on this channel instance can be entered simultaneously with this one.

When you write a channel exit program, consider whether it will run on a channel instance that might involve sharing conversations. If the channel instance might involve sharing conversations, consider the effect on other instances of the channel exit of changing MQCD fields; all MQCD fields have common values across all the sharing conversations. After the channel instance is established, if exit programs try to alter MQCD fields they might encounter problems because other instances of exit programs running on the channel instance could be attempting to alter the same fields at the same time. If this situation could arise with your exit programs, you must serialize access to the MQCD in your exit code.

If you are working with a channel which is defined to share conversations, but you do not want sharing to occur on a particular channel instance, set the MQCD value of SharingConversations to 1 or 0 when you initialize a channel exit on the first conversation on the channel instance. See SharingConversations for an explanation of the values of SharingConversations.

Example

Sharing conversations is enabled.

You are using a client-connection channel definition which specifies an exit program.

The first time that this channel starts, the exit program alters some of the MQCD parameters when it is initialized. These are acted on by the channel, so the definition that the channel is running with is now different from the one that was originally supplied. The MQCXP SharingConversations parameter is set to TRUE.

The next time that the application connects using this channel, the conversation runs on the channel instance which was started previously, because it has the same original channel definition. The channel instance the application connects to the second time is the same instance as the first time it connected. Consequently, it uses the definitions that have been altered by the exit program. When the exit program is initialized for the second conversation, although it can alter MQCD fields, they are not acted on by the channel. These same characteristics apply to any subsequent conversations which share the channel instance.

Using MQCONNX:

You can use the MQCONNX call to specify a channel definition (MQCD) structure in the MQCNO structure.

This allows the calling client application to specify the definition of the client-connection channel at run time. For more information, see Using the MQCNO structure on an MQCONNX call. When you use MQCONNX, the call issued at the server depends on the server level and listener configuration.

When you use MQCONNX from a client, the following options are ignored:

- MQCNO_STANDARD_BINDING
- MQCNO_FASTPATH_BINDING

The MQCD structure you can use depends on the MQCD version number you are using. For information on MQCD versions (MQCD_VERSION), see MQCD Version. You can use the MQCD structure, for example, to pass channel-exit programs to the server. If you are using MQCD Version 3 or later, you can use the structure to pass an array of exits to the server. You can use this function to perform more than one operation on the same message, such as encryption and compression, by adding an exit for each operation, rather than modifying an existing exit. If you do not specify an array in the MQCD structure, the single exit fields will be checked. For more information on channel-exit programs, see “Channel-exit programs for messaging channels” on page 341.

Shared connection handles on MQCONNX:

You can share handles between different threads within the same process, using shared connection handles.

When you specify a shared connection handle, the connection handle returned from the MQCONNX call can be passed in subsequent MQI calls on any thread in the process.

Note: You can use a shared connection handle on an IBM MQ MQI client to connect to a server queue manager that does not support shared connection handles.

For more information, see “Using MQCONNX.”

Building applications for IBM MQ MQI clients

Applications can be built and run in IBM MQ MQI client environment. The application must be built and linked to the IBM MQ MQI client used. The way in which applications are built and linked varies according to the platform and programming language used.

If an application is to run in a client environment, you can write it in the languages shown in the following table:

Table 34. Programming languages supported in client environments






Client platform	C	C++	COBOL	pTAL	RPG	Visual Basic
 AIX	Yes	Yes	Yes			
 NSS Client HP Integrity NonStop Server	Yes		Yes	Yes		
 HP-UX	Yes	Yes	Yes			
 IBM i	Yes		Yes		Yes	
 Linux	Yes	Yes	Yes			

Table 34. Programming languages supported in client environments (continued)

Client platform	C	C++	COBOL	pTAL	RPG	Visual Basic
Solaris Solaris	Yes	Yes	Yes			
Windows Windows	Yes	Yes	Yes			Yes

Linking C applications with the IBM MQ MQI client code: **Multi**

Having written your IBM MQ application that you want to run on the IBM MQ MQI client, you must link it to the IBM MQ MQI client code.

You can link your application to the IBM MQ MQI client code in two ways:

1. Directly, by connecting your application to a queue manager, in which case the queue manager must be on the same machine as your application.
2. To a client library file, which gives you access to queue managers on the same or on a different machine.

IBM MQ provides a client library file for each environment:

AIX **AIX**

libmqic.a library for non-threaded applications, or libmqic_r.a library for threaded applications.

HP-UX **HP-UX**

libmqic.sl library for non-threaded applications, or libmqic_r.sl library for threaded applications.

Linux **Linux**

libmqic.so library for non-threaded applications, or libmqic_r.so library for threaded applications.

IBM i **IBM i**

Bind client application with LIBMQIC client service program for non-threaded applications, or LIBMQIC_R service program for threaded applications.

Solaris **Solaris**

libmqic.so.

If you want to use the programs on a machine that has only the IBM MQ MQI client for Solaris installed, you must recompile the programs to link them with the client library:

```
$ /opt/SUNWsprow/bin/cc -o prog_name prog_name.c -mt -lmqic \
-lsocket -lc -lnsl -ldl
```

The parameters must be entered in the correct order, as shown.

Windows **Windows**

MQIC32.LIB.

Linking C++ applications with the IBM MQ MQI client code: **ULW**

You can write applications to run on the client in C++. Build methods vary according to the environment.

For information about how to link your C++ applications, see Building IBM MQ C++ programs.

For full details of all aspects of using C++, see Using C++

Linking COBOL applications with the IBM MQ MQI client code: **Multi**

Having written a COBOL application that you want to run on the IBM MQ MQI client, you must link it with an appropriate library.

IBM MQ provides a client library file for each environment:

AIX **AIX**

Link your non-threaded COBOL application with the library libmqicb.a or threaded COBOL application with libmqicb_r.a.

HP-UX **HP-UX**

Link your non-threaded COBOL application with the library libmqicb.sl or threaded COBOL application with libmqicb_r.sl.

HP-UX **Linux**

Link your non-threaded COBOL application with the library libmqicb.so or threaded COBOL application with libmqicb_r.so.

IBM i **IBM i**

Bind COBOL client application with AMQCSTUB service program for non-threaded applications, or AMQCSTUB_R service program for threaded applications.

Solaris **Solaris**

Link your non-threaded COBOL application with the library libmqicb.so or threaded COBOL application with libmqicb_r.so.

Windows **Windows**

Link your application code with the MQICCB library for 32-bit COBOL. The IBM MQ MQI client for Windows does not support 16-bit COBOL.

Linking Visual Basic applications with the IBM MQ MQI client code: **Windows**

You can link Microsoft Visual Basic applications with the IBM MQ MQI client code on Windows.

V 9.0.0 From Version 9.0, IBM MQ support for Microsoft Visual Basic 6.0 is deprecated. IBM MQ classes for .NET are the recommended replacement technology. For more information, see Developing .NET applications.

Link your Visual Basic application with the following include files:

CMQB.bas
MQI

CMQBB.bas
MQAI

CMQCFB.bas
PCF commands

CMQXB.bas

Channels

Set `motype=2` for the client in the Visual Basic compiler, to ensure the correct automatic selection of the client dll:

MQIC32.dll

Windows 7, Windows 8, Windows 2008, and Windows 2012

Related concepts:

“Coding in Visual Basic” on page 450

Information to consider when coding IBM MQ programs in Microsoft Visual Basic. Visual Basic is supported only on Windows.

“Preparing Visual Basic programs in Windows” on page 417

Information to consider when using Microsoft Visual Basic programs on Windows.

Running applications in the IBM MQ MQI client environment

You can run an IBM MQ application both in a full IBM MQ environment and in an IBM MQ MQI client environment without changing your code, provided that certain conditions are met.

These conditions are that:

- The application does not need to connect to more than one queue manager concurrently.
- The queue manager name is not prefixed with an asterisk (*) on an MQCONN or MQCONNX call.
- The application does not need to use any of the exceptions listed in What applications run on an IBM MQ MQI client?

Note: The libraries that you use at link-edit time determine the environment in which your application must run.

When working in the IBM MQ MQI client environment, remember that:






- Each application running in the IBM MQ MQI client environment has its own connections to servers. An application establishes one connection to a server each time it issues an MQCONN or MQCONNX call.
- An application sends and gets messages synchronously. This implies a wait between the time the call is issued at the client and the return of a completion code and reason code across the network.
- All data conversion is done by the server, but see also MQCCSID for information about overriding the machine's configured CCSID.

Connecting IBM MQ MQI client applications to queue managers:

An application running in an IBM MQ MQI client environment can connect to a queue manager in various ways. You can use environment variables, the MQCNO structure, or a client definition table.

When an application running in an IBM MQ client environment issues an MQCONN or MQCONNX call, the client identifies how it is to make the connection. When an MQCONNX call is issued by an application on an IBM MQ client, the MQI client library searches for the client channel information in the following order:

1. Using the contents of the *ClientConnOffset* or *ClientConnPtr* fields of the MQCNO structure (if supplied). These fields identify the channel definition structure (MQCD) to be used as the definition of the client connection channel. Connection details can be overridden by using a pre-connect exit. For more information, see “Referencing connection definitions using a pre-connect exit from a repository” on page 374.
2. If the MQSERVER environment variable is set, the channel that it defines is used.
3. If an `mqclient.ini` file is defined and contains a `ServerConnectionParms`, the channel that it defines is used. For more information, see Configuring a client using a configuration file and CHANNELS stanza of the client configuration file.


4. If the MQCHLLIB and MQCHLTAB environment variables are set, the client channel definition table that they point to is used.  Alternatively, from IBM MQ Version 9.0, the MQCCDTURL environment variable provides the equivalent capability to setting a combination of the MQCHLLIB and MQCHLTAB environment variables. If MQCCDTURL is set, the client channel definition table that it points to is used. For more information, see Web addressable access to the client channel definition table.
5. If an mqclient.ini file is defined and contains ChannelDefinitionDirectory and ChannelDefinitionFile attributes, these attributes are used to locate the client channel definition table. For more information, see Configuring a client using a configuration file and CHANNELS stanza of the client configuration file.
6. Finally, if the environment variables are not set, the client searches for a client channel definition table with a path and name that are established from the DefaultPrefix in the mqs.ini file. If the search for a client channel definition table fails, the client uses the following paths:
 -   On UNIX and Linux: /var/mqm/AMQCLCHL.TAB
 -  On Windows: C:\Program Files\IBM\MQ\amqclchl.tab
 -  On IBM i: /QIBM/UserData/mqm/@ipcc
 - On IBM MQ Appliance: QMname_AMQCLCHL.TAB. They appear under the mqbackup:// URI.

The first of the options described in the previous list (using the *ClientConnOffset* or *ClientConnPtr* fields of MQCNO) is supported only by the MQCONN call. If the application is using MQCONN rather than MQCONNX, the channel information is searched for in the remaining five ways in the order shown in the list. If the client fails to find the channel information, the MQCONN or MQCONNX call fails.

The channel name (for the client connection) must match the server-connection channel name defined on the server for the MQCONN or MQCONNX call to succeed.

Related information:

Client channel definition table

 Web addressable access to the client channel definition table

MQSERVER

MQCHLLIB

MQCHLTAB

 MQCCDTURL

Configuring connections between the server and client

Connecting client applications to queue managers using environment variables:

Client channel information can be supplied to an application running in a client environment by environment variables.

An application running in an IBM MQ MQI client environment can connect to a queue manager by using the following environment variables:

MQSERVER


The 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.

MQCHLLIB

The MQCHLLIB environment variable 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.

MQCHLTAB

The MQCHLTAB environment variable specifies the name of the file containing the client channel definition table (CCDT).

 From IBM MQ Version 9.0, the MQCCDTURL environment variable provides the equivalent capability to setting a combination of the MQCHLLIB and MQCHLTAB environment variables. MQCCDTURL allows you to provide a file, ftp, or http URL as a single value from which a client channel definition table can be obtained. For more information, see [Web addressable access to the client channel definition table](#).

Connecting client applications to queue managers using the MQCNO structure:

You can specify the definition of the channel in a channel definition structure (MQCD), which is supplied using the MQCNO structure of the MQCONN call.

For more information see [Using the MQCNO structure on an MQCONN call](#).

Connecting client applications to queue managers using a client channel definition table:

If you use the MQSC DEFINE CHANNEL command, the details you provide are placed in the client channel definition table (ccdt). The contents of the **QMGrName** parameter of the MQCONN or MQCONN call determines which queue manager the client connects to.

This file is accessed by the client to determine the channel an application will use. Where there is more than one suitable channel definition, the choice of channel is influenced by the client channel weight (CLNTWGHT) and connection affinity (AFFINITY) channel attributes.

Using 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 MQCONN, and to reopen objects. Automatic client reconnection is suitable for many, but not all client applications.

For more information, see [Automatic client reconnection](#).

Role of the client channel definition table:

The client channel definition table (CCDT) contains definitions of client connection channels. It is particularly useful if your client applications might need to connect to a number of alternative queue managers.

The client channel definition table is created when you define a queue manager. The same file can be used by more than one IBM MQ client.

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.

V 9.0.0 From IBM MQ Version 9.0, the CCDT can be hosted in a central location that is accessible through a URI, removing the need to individually update the CCDT for each deployed client.

Related information:

Client channel definition table

V 9.0.0 Web addressable access to the client channel definition table

Accessing client-connection channel definitions

Queue manager groups in the CCDT:

You can define a set of connections in the client channel definition table (CCDT) as a *queue manager group*. You can connect an application to a queue manager that is part of a queue manager group. This can be done by prefixing the queue manager name on an MQCONN or MQCONNX call with an asterisk.

You might choose to define connections to more than one server machine because:

- You want to connect a client to any one of a set of queue managers that is running, to improve availability.
- You want to reconnect a client to the same queue manager it connected to successfully last time, but connect to a different queue manager if the connection fails.
- You want to be able to retry a client connection to a different queue manager if the connection fails, by issuing the MQCONN in the client program again.
- You want to automatically reconnect a client connection to another queue manager if the connection fails, without writing any client code.
- You want to automatically reconnect a client connection to a different instance of a multi-instance queue manager if a standby instance takes over, without writing any client code.
- You want to balance your client connections across a number of queue managers, with more clients connecting to some queue managers than others.
- You want to spread the reconnection of many client connections over multiple queue managers and over time, in case the high volume of connections causes a failure.
- You want to be able to move your queue managers without changing any client application code.
- You want to write client application programs that do not need to know queue manager names.

It is not always appropriate to connect to different queue managers. An extended transactional client or a Java client in WebSphere Application Server, for example, might need to connect to a predictable queue manager instance. Automatic client reconnect is not supported by IBM MQ classes for Java.

A queue manager group is a set of connections defined in the client channel definition table (CCDT). The set is defined by its members having the same value of the **QMNAME** attribute in their channel definitions.

Figure 50 is a graphical representation of a client connection table, showing three queue manager groups, two named queue manager groups written in the CCDT as **QMNAME** (QM1) and **QMNAME** (QMGrp1), and one blank or default group written as **QMNAME** (' ').

1. Queue manager group QM1 has three client connection channels, connecting it to queue managers QM1 and QM2. QM1 might be a multi-instance queue manager located on two different servers.
2. The default queue manager group has six client connection channels connecting it to all the queue managers.
3. QMGrp1 has client connection channels to two queue managers, QM4 and QM5.

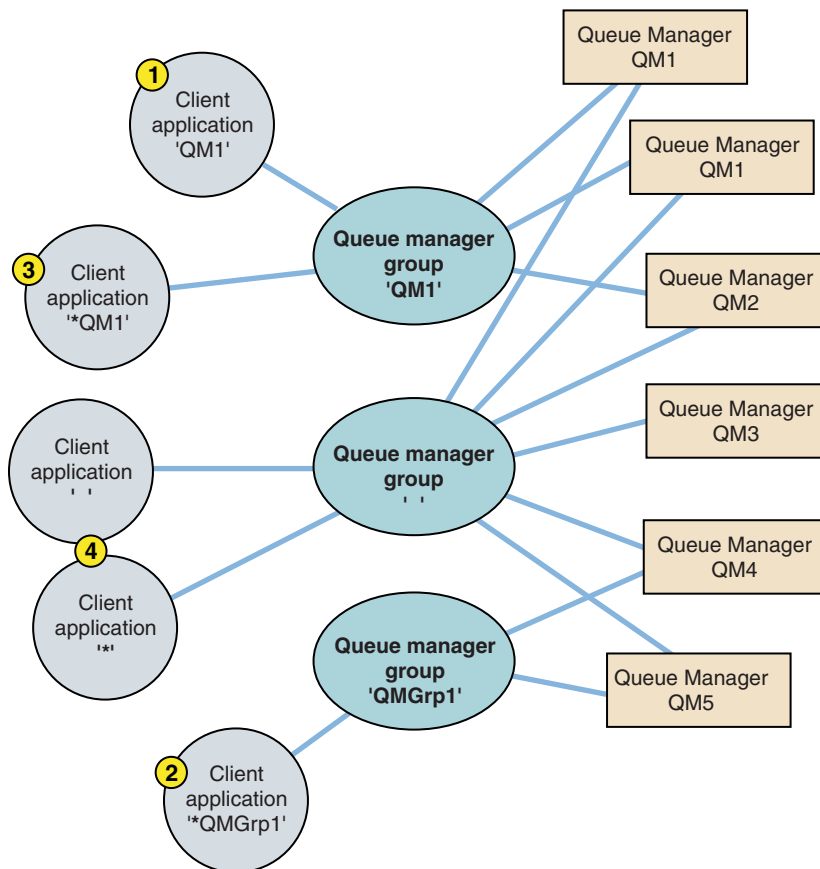


Figure 50. Queue manager groups

Four examples of using this client connection table are described with the help of the numbered client applications in Figure 50.

1. In the first example, the client application passes a queue manager name, QM1, as the **QmgrName** parameter to its MQCONN or MQCONNX MQI call. The IBM MQ client code selects the matching queue manager group, QM1. The group contains three connection channels, and the IBM MQ MQI client tries to connect to QM1 using each of these channels in turn until it finds an IBM MQ listener for the connection attached to a running queue manager called QM1.

The order of connection attempts depends on the value of the client connection AFFINITY attribute and the client channel weightings. Within these constraints, the order of connection attempts is randomized, both over the three possible connections, and over time, in order to spread out the load of making connections.

The MQCONN or MQCONNX call issued by the client application succeeds when a connection is established to a running instance of QM1.

2. In the second example, the client application passes a queue manager name prefixed with an asterisk, *QMGrp1 as the **QmgrName** parameter to its MQCONN or MQCONNX MQI call. The IBM MQ client selects the matching queue manager group, QMGrp1. This group contains two client connection channels, and the IBM MQ MQI client tries to connect to *any* queue manager using each channel in turn. In this example, the IBM MQ MQI client needs to make a successful connection; the name of the queue manager that it connects to does not matter.

The rule for the order of making connection attempts is the same as before. The only difference is that by prefixing the queue manager name with an asterisk, the client indicates that the name of the queue manager is not relevant.

The MQCONN or MQCONNX call issued by the client application succeeds when a connection is established to a running instance of any queue manager connected to by the channels in the QMGrp1 queue manager group.

3. The third example is essentially the same as the second because the **QmgrName** parameter is prefixed by an asterisk, *QM1. The example illustrates that you cannot determine which queue manager a client channel connection is going to connect to by inspecting the QMNAME attribute in one channel definition by itself. The fact that the **QMNAME** attribute of the channel definition is QM1, is not sufficient to require a connection is made to a queue manager called QM1. If your client application prefixes its **QmgrName** parameter with an asterisk then any queue manager is a possible connection target.

In this case the MQCONN or MQCONNX calls issued by the client application succeed when a connection is established to a running instance of either QM1 or QM2.

4. The fourth example illustrates use of the default group. In this case the client application passes an asterisk, '*', or blank ' ', as the **QmgrName** parameter to its MQCONN or MQCONNX MQI call. By convention in the client channel definition, a blank **QMNAME** attribute signifies the default queue manager group and either a blank or asterisk **QmgrName** parameter matches a blank **QMNAME** attribute.

In this example the default queue manager group has client channel connections to all the queue managers. By selecting the default queue manager group the application might be connected to any queue manager in the group.

The MQCONN or MQCONNX call issued by the client application succeeds when a connection is established to a running instance of any queue manager.

Note: The default group is different from a default queue manager, although an application uses a blank **QmgrName** parameter to connect to either the default queue manager group or to the default queue manager. The concept of a default queue manager group is only relevant to a client application, and a default queue manager to a server application.

Define your client connection channels on one queue manager only, including those channels that connect to a second or third queue manager. Do not define them on two queue managers and then try to merge the two client channel definition tables. Only one client channel definition table can be accessed by the client.

Examples

Look again at the list of reasons for using queue manager groups at the beginning of the topic. How does using a queue manager group provide those capabilities?

Connect to any one of a set of queue managers.

Define a queue manager group with connections to all the queue managers in the set, and connect to the group using the **QmgrName** parameter prefixed by an asterisk.

Reconnect to the same queue manager, but connect to a different one, if the queue manager connected to last time is unavailable.

Define a queue manager group as before but set the attribute, **AFFINITY** (PREFERRED) on each client channel definition.

Retry a connection to another queue manager if a connection fails.

Connect to a queue manager group, and reissue the MQCONN or MQCONNX MQI call if the connection is broken or the queue manager fails.

Automatically reconnect to another queue manager if a connection fails.

Connect to a queue manager group using the MQCONNX **MQCNO** option MQCNO_RECONNECT.

Automatically reconnect to a different instance of a multi-instance queue manager.

Do the same as the preceding example. In this case, if you want to restrict the queue manager group to connect to the instances of a particular multi-instance queue manager, define the group with connections to only the multi-instance queue manager instances.

You can also ask the client application to issue its MQCONN or MQCONNX MQI call with no asterisk prefixed to the **QmgrName** parameter. That way the client application can only connect to the named queue manager. Finally, you can set the **MQCNO** option to MQCNO_RECONNECT_Q_MGR. This option accepts reconnections to the same queue manager that was previously connected. You can also use this value to restrict reconnections to the same instance of a normal queue manager.

Balance client connections across queue managers, with more clients connected to some queue managers than others.

Define a queue manager group, and set the **CLNTWGHT** attribute on each client channel definition to distribute the connections unevenly.

Spread the client reconnection load unevenly, and spread it over time, after a connection or queue manager failure.


Do the same as the preceding example. The IBM MQ MQI client randomizes reconnections across queue managers and spreads the reconnections over time.

Move your queue managers without changing any client code.

The CCDT isolates your client application from the location of the queue manager. The CCDT is a data file which can be defined at the client, read from a shared location or fetched from a web-server. For more information, see Client channel definition table.

Write a client application that does not know queue manager names.

Use queue manager group names and establish a naming convention for queue manager group names that is relevant to your client applications in your organization, and reflects the architecture of your solutions rather than the naming of queue managers.

Connecting to queue-sharing groups: 

You can connect your application to a queue manager that is part of a queue-sharing group. This can be done by using the queue-sharing group name instead of the queue manager name on the MQCONN or MQCONNX call.

Queue-sharing groups have a name of up to four characters. The name must be unique in your network, and must be different from any queue manager names.

The client channel definition should use the queue-sharing group generic interface to connect to an available queue manager in the group. For more information, see Connecting a client to a queue-sharing group. A check is made to ensure that the queue manager the listener connects to is a member of the queue-sharing group.

For more information on shared queues, see Shared queues and queue-sharing groups.

Examples of channel weighting and affinity:

These examples illustrate how client-connection channels are selected when non-zero `ClientChannelWeights` are used.

The `ClientChannelWeight` and `ConnectionAffinity` channel attributes control how client-connection channels are selected when more than one suitable channel is available for a connection. These channels are configured to connect to different queue managers in order to provide higher availability, workload balancing, or both. MQCONN calls that could result in a connection to one of several queue managers must prefix the queue manager name with an asterisk as described in: Examples of MQCONN calls: Example 1. Queue manager name includes an asterisk (*).

Applicable candidate channels for a connection are those where the `QMNAME` attribute matches the queue manager name specified in the MQCONN call. If all applicable channels for a connection have a `ClientChannelWeight` of zero (the default) then they are selected in alphabetical order as in the example: Examples of MQCONN calls: Example 1. Queue manager name includes an asterisk (*).

The following examples illustrate what happens when non-zero `ClientChannelWeights` are used. Note that, since this feature involves pseudo-random channel selection, the examples show a sequence of actions that might happen rather than what definitely will.

Example 1. Selecting channels when ConnectionAffinity is set to PREFERRED:

This example illustrates how an IBM MQ MQI client selects a channel from a CCDT, where the `ConnectionAffinity` is set to `PREFERRED`.

In this example, a number of client machines use a Client Channel Definition Table (CCDT) provided by a queue manager. The CCDT includes client connection channels with the following attributes (shown using the syntax of the `DEFINE CHANNEL` command):

```
CHANNEL(A) QMNAME(DEV) CONNAME(devqm.it.company.example)
CHANNEL(B) QMNAME(CORE) CONNAME(core1.ops.company.example) CLNTWGHT(5) +
AFFINITY(PREFERRED)
CHANNEL(C) QMNAME(CORE) CONNAME(core2.ops.company.example) CLNTWGHT(3) +
AFFINITY(PREFERRED)
CHANNEL(D) QMNAME(CORE) CONNAME(core3.ops.company.example) CLNTWGHT(2) +
AFFINITY(PREFERRED)
```

The application issues `MQCONN(*CORE)`

Channel A is not a candidate for this connection, because the `QMNAME` attribute does not match. Channels B, C and D are identified as candidates, and are placed in an order of preference based on their weighting. In this example the order might be B, C, D. The client attempts to connect to the queue manager at `core2.ops.company.example`. The name of the queue manager at that address is not checked, because the MQCONN call included an asterisk in the queue manager name.

It is important to note that, with `AFFINITY(PREFERRED)`, each time this particular client machine connects it will place the channels in the same initial order of preference. This applies even when the connections are from different processes or at different times.

In this example, the queue manager at `core2.ops.company.example` cannot be reached. The client attempts to connect to `core1.ops.company.example` because channel B is next in the order of preference. In addition, channel C is demoted to become the least preferred.

A second `MQCONN(*CORE)` call is issued by the same application. Channel C was demoted by the previous connection, so the most preferred channel is now B. This connection is made to `core1.ops.company.example`.

A second machine sharing the same Client Channel Definition Table might place the channels in a different initial order of preference. For example, D, B, C. Under normal circumstances, with all channels working, applications on this machine are connected to core3.ops.company.example while those on the first machine are connected to core2.ops.company.example. This allows workload balancing of large numbers of clients across multiple queue managers while allowing each individual client to connect to the same queue manager if it is available.

Example 2. Selecting channels when ConnectionAffinity is set to NONE:

This example illustrates how an IBM MQ MQI client selects a channel from a CCDT, where the ConnectionAffinity is set to NONE.

In this example, a number of clients use a Client Channel Definition Table (CCDT) provided by a queue manager. The CCDT includes client connection channels with the following attributes (shown using the syntax of the DEFINE CHANNEL command):

```
CHANNEL(A) QMNAME(DEV) CONNAME(devqm.it.company.example)
CHANNEL(B) QMNAME(CORE) CONNAME(core1.ops.company.example) CLNTWGHT(5) +
AFFINITY(NONE)
CHANNEL(C) QMNAME(CORE) CONNAME(core2.ops.company.example) CLNTWGHT(3) +
AFFINITY(NONE)
CHANNEL(D) QMNAME(CORE) CONNAME(core3.ops.company.example) CLNTWGHT(2) +
AFFINITY(NONE)
```

The application issues MQCONN(*CORE). As in the previous example, channel A is not considered because the QMNAME does not match. Channel B, C, or D are selected based on their weighting, with probabilities of 50%, 30%, or 20%. In this example, channel B might be selected. There is no persistent order of preference created.

A second MQCONN(*CORE) call is made. Again, one of the three applicable channels is selected, with the same probabilities. In this example, channel C is chosen. However, core2.ops.company.example does not respond, so another choice is made between the remaining candidate channels. Channel B is selected and the application is connected to core1.ops.company.example.

With AFFINITY(NONE), each MQCONN call is independent of any other. Therefore when this example application makes a third MQCONN(*CORE), it might once more attempt to connect through the broken channel C, before choosing one of B or D.

Examples of MQCONN calls:

Examples of using MQCONN to connect to a specific queue manager, or to one of a group of queue managers.

In each of the following examples, the network is the same; there is a connection defined to two servers from the same IBM MQ MQI client. (In these examples, the MQCONN call could be used instead of the MQCONN call.)

There are two queue managers running on the server machines, one named SALE and the other named SALE_BACKUP.

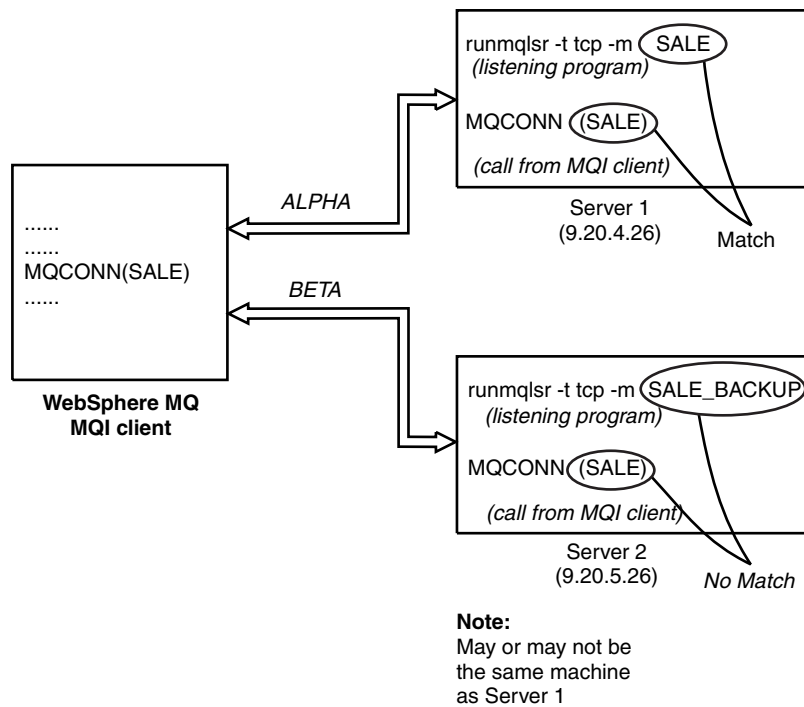


Figure 51. MQCONN example

The definitions for the channels in these examples are:

SALE definitions:

```

DEFINE CHANNEL(ALPHA) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
DESCR('Server connection to IBM MQ MQI client')

DEFINE CHANNEL(ALPHA) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNNAME(9.20.4.26) DESCR('IBM MQ MQI client connection to server 1') +
QMNAME(SALE)

DEFINE CHANNEL(BETA) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNNAME(9.20.5.26) DESCR('IBM MQ MQI client connection to server 2') +
QMNAME(SALE)

```

SALE_BACKUP definition:

```

DEFINE CHANNEL(BETA) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
DESCR('Server connection to IBM MQ MQI client')

```

The client channel definitions can be summarized as follows:

Name	CHLTYPE	TRPTYPE	CONNAME	QMNAME
ALPHA	CLNTCONN	TCP	9.20.4.26	SALE
BETA	CLNTCONN	TCP	9.20.5.26	SALE

What the MQCONN examples demonstrate:

The examples demonstrate the use of multiple queue managers as a backup system.

Suppose the communication link to Server 1 is temporarily broken. The use of multiple queue managers as a backup system is demonstrated.

Each example covers a different MQCONN call and gives an explanation of what happens in the specific example presented, by applying the following rules:

1. The client channel definition table (CCDT) is scanned in alphabetical channel name order for a queue manager name (QMNAME field) corresponding to the one given in the MQCONN call.
2. If a match is found, the channel definition is used.
3. An attempt is made to start the channel to the machine identified by the connection name (CONNAME). If this is successful, the application continues. It requires:
 - A listener to be running on the server.
 - The listener to be connected to the same queue manager as the one the client wants to connect to (if specified).
4. If the attempt to start the channel fails and there is more than one entry in the client channel definition table (in this example there are two entries), the file is searched for a further match. If a match is found, processing continues at step 1.
5. If no match is found, or there are no more entries in the client channel definition table and the channel has failed to start, the application is unable to connect. An appropriate reason code and completion code are returned in the MQCONN call. The application can take action based on the reason and completion codes returned.

Example 1. Queue manager name includes an asterisk ():*

In this example the application is not concerned about which queue manager it connects to. The application issues an MQCONN call for a queue manager name including an asterisk. A suitable channel is chosen.

The application issues:

```
MQCONN (*SALE)
```

Following the rules, this is what happens in this instance:

1. The client channel definition table (CCDT) is scanned for the queue manager name SALE, matching with the application MQCONN call.
2. Channel definitions for ALPHA and BETA are found.
3. If one channel has a CLNTWGHT value of 0, this channel is selected. If both have a CLNTWGHT value of 0, channel ALPHA is selected because it is first in alphabetical sequence. If both channels have a non-zero CLNTWGHT value, one channel is randomly selected, based on its weighting.
4. An attempt to start the channel is made.
5. If channel BETA was selected, the attempt to start it is successful.
6. If channel ALPHA was selected, the attempt to start it is NOT successful because the communication link is broken. The following steps then apply:
 - a. The only other channel for the queue manager name SALE is BETA.

- b. An attempt to start this channel is made - this is successful.
7. A check to see that a listener is running shows that there is one running. It is not connected to the SALE queue manager, but because the MQI call parameter has an asterisk (*) included in it, no check is made. The application is connected to the SALE_BACKUP queue manager and continues processing.

Example 2. Queue manager name specified:

In this example the application must connect to a particular queue manager. The application issues an MQCONN call for that queue manager name. A suitable channel is chosen.

The application requires a connection to a specific queue manager, named SALE, as seen in the MQI call:
MQCONN (SALE)

Following the rules, this is what happens in this instance:

1. The client channel definition table (CCDT) is scanned in alphabetical channel name sequence, for the queue manager name SALE, matching with the application MQCONN call.
2. The first channel definition found to match is ALPHA.
3. An attempt to start the channel is made - this is not successful because the communication link is broken.
4. The client channel definition table is again scanned for the queue manager name SALE and the channel name BETA is found.
5. An attempt to start the channel is made - this is successful.
6. A check to see that a listener is running shows that there is one running, but it is not connected to the SALE queue manager.
7. There are no further entries in the client channel definition table. The application cannot continue and receives return code MQRC_Q_MGR_NOT_AVAILABLE.

Example 3. Queue manager name is blank or an asterisk ():*

In this example the application is not concerned about which queue manager it connects to. The application issues an MQCONN specifying a blank queue manager name or an asterisk. A suitable channel is chosen.

This is treated in the same way as "Example 1. Queue manager name includes an asterisk (*)" on page 304.

Note: If this application were running in an environment other than an IBM MQ MQI client, and the name was blank, it would be attempting to connect to the default queue manager. This is not the case when it is run from a client environment; the queue manager accessed is the one associated with the listener to which the channel connects.

The application issues:

```
MQCONN ("")
```

or

```
MQCONN (*)
```

Following the rules, this is what happens in this instance:

1. The client channel definition table (CCDT) is scanned in alphabetical channel name sequence, for a queue manager name that is blank, matching with the application MQCONN call.
2. The entry for the channel name ALPHA has a queue manager name in the definition of SALE. This does not match the MQCONN call parameter, which requires the queue manager name to be blank.
3. The next entry is for the channel name BETA.

4. The queue manager name in the definition is SALE. Once again, this does not match the MQCONN call parameter, which requires the queue manager name to be blank.
5. There are no further entries in the client channel definition table. The application cannot continue and receives return code MQRC_Q_MGR_NOT_AVAILABLE.

Triggering in the client environment:

Messages sent by IBM MQ applications running on IBM MQ MQI clients contribute to triggering in exactly the same way as any other messages, and they can be used to trigger programs on both the server and the client.

Triggering is explained in detail in the Triggering channels.

The trigger monitor and the application to be started must be on the same system.

The default characteristics of the triggered queue are the same as those in the server environment. In particular, if no MQPMO sync point control options are specified in a client application putting messages to a triggered queue that is local to a z/OS queue manager, the messages are put within a unit of work. If the triggering condition is then met, the trigger message is put on the initiation queue within the same unit of work and cannot be retrieved by the trigger monitor until the unit of work ends. The process that is to be triggered is not started until the unit of work ends.

Process definition:


You must define the process definition on the server, because this is associated with the queue that has triggering set on.

The process object defines what is to be triggered. If the client and server are not running on the same platform, any processes started by the trigger monitor must define *ApplType*, otherwise the server takes its default definitions (that is, the type of application that is normally associated with the server machine) and causes a failure.



For example, if the trigger monitor is running on a Windows client and wants to send a request to a server on another operating system, MQAT_WINDOWS_NT must be defined otherwise the other operating system uses its default definitions and the process fails.

Trigger monitor:

The trigger monitor provided by non-z/OS IBM MQ products runs in the client environments for

 IBM i, UNIX, Linux, and Windows systems.

To run the trigger monitor, issue one of these commands:

-  On IBM i:
`CALL PGM(QMQM/RUNMQTMC) PARM('-m' QmgrName '-q' InitQ)`
-  On Windows, UNIX and Linux platforms:
`runmqtmc [-m QMgrName] [-q InitQ]`

The default initiation queue is SYSTEM.DEFAULT.INITIATION.QUEUE on the default queue manager. The initiation queue is where the trigger monitor looks for trigger messages. It then calls programs for the appropriate trigger messages. This trigger monitor supports the default application type and is the same as `runmqtrm` except that it links the client libraries.

The command string, built by the trigger monitor, is as follows:

1. The *ApplicId* from the relevant process definition. *ApplicId* is the name of the program to run, as it would be entered on the command line.
2. The MQTMC2 structure, enclosed in quotation marks, obtained from the initiation queue. A command string is started that has this string, exactly as provided, in quotation marks in order that the system command accepts it as one parameter.
3. The *EnvrData* from the relevant process definition.

The trigger monitor does not look to see if there is another message on the initiation queue until the completion of the application it has started. If the application has much processing to do, the trigger monitor might not keep up with the number of trigger messages arriving. There are two ways to deal with this situation:

1. Have more trigger monitors running
If you choose to have more trigger monitors running, you can control the maximum number of applications that can run at any one time.
2. Run the started applications in the background
If you choose to run applications in the background, IBM MQ imposes no restriction on the number of applications that can run.

To run the started application in the background on UNIX and Linux systems, you must put an & (ampersand) at the end of the *EnvrData* of the process definition.

CICS applications (non-z/OS):

A non-z/OS CICS application program that issues an MQCONN or MQCONNX call must be defined to CEDA as RESIDENT. If you relink a CICS server application as a client, you risk losing sync point support.

A non-z/OS CICS application program that issues an MQCONN or MQCONNX call must be defined to CEDA as RESIDENT. To make the resident code as small as possible, you can link to a separate program to issue the MQCONN or MQCONNX call.

If the MQSERVER environment variable is used to define the client connection, it must be specified in the CICSENV.CMD file.

IBM MQ applications can be run in an IBM MQ server environment or on an IBM MQ client without changing code. However, in an IBM MQ server environment, CICS can act as sync point coordinator, and you use EXEC CICS SYNCPOINT and EXEC CICS SYNCPOINT ROLLBACK rather than **MQCMIT** and **MQBACK**. If a CICS application is simply relinked as a client, sync point support is lost. **MQCMIT** and **MQBACK** must be used for the application running on an IBM MQ MQI client.

Preparing and running CICS and Tuxedo applications

UNIX

Windows

To run CICS and Tuxedo applications as client applications, you use different libraries from those you use with server applications. The user ID under which the application runs is also different.

To prepare CICS and Tuxedo applications to run as IBM MQ MQI client applications, follow the instructions in the *Configuring an extended transactional client*.

Note, however, that the information that deals specifically with preparing CICS and Tuxedo applications, including the sample programs supplied with IBM MQ, assumes that you are preparing applications to run on an IBM MQ server system. As a result, the information refers only to IBM MQ libraries that are intended for use on a server system. When you are preparing your client applications, you must do the following things:

- Use the appropriate client system library for the language bindings that your application uses. For example:
 - **UNIX** For applications written in C on UNIX, use the library `libmqic` instead of `libmqm`.
 - **Windows** On Windows systems, use the library `mqic.lib` instead of `mqm.lib`.
- Instead of the server system libraries shown in Table 35 and Table 36, use the equivalent client system libraries. If a server system library is not listed in these tables, use the same library on a client system.

Table 35. Client system libraries on UNIX

Library for an IBM MQ server system	Equivalent library to use on an IBM MQ client system
<code>libmqmxa</code>	<code>libmqcxa</code>

Table 36. Client system libraries on Windows systems

Library for an IBM MQ server system	Equivalent library to use on an IBM MQ client system
<code>mqmxa.lib</code>	<code>mqcxa.lib</code>
<code>mqmtux.lib</code>	<code>mqcxa.lib</code>
<code>mqmenc.lib</code>	<code>mqcxa.lib</code>
<code>mqmcics4.lib</code>	<code>mqccics4.lib</code>

The user ID used by a client application

When you run an IBM MQ server application under CICS, it normally switches from the CICS user to the user ID of the transaction. However, when you run an IBM MQ MQI client application under CICS, it retains the CICS privileged authority.

CICS and Tuxedo sample programs: UNIX Windows

CICS and Tuxedo sample programs for use on UNIX and Windows systems.

Table 37 lists the CICS and Tuxedo sample programs that are supplied for use on UNIX client systems. Table 38 lists the equivalent information for Windows client systems. The tables also list the files that are used for preparing and running the programs. For a description of the sample programs, see “The CICS transaction sample” on page 481 and “Using the TUXEDO samples on UNIX and Windows” on page 524.

Table 37. Sample programs for UNIX client systems

Description	Source	Executable module
CICS program	amqscic0.ccs	amqscicc
Header file for the CICS program	amqscih0.h	-
Tuxedo client program to put messages	amqstxpx.c	-
Tuxedo client program to get messages	amqstxgx.c	-
Tuxedo server program for the two client programs	amqstxsx.c	-
UBBCONFIG file for the Tuxedo programs	ubbstxcx.cfg	-
Field table file for the Tuxedo programs	amqstxvx.flds	-
View description file for the Tuxedo programs	amqstxvx.v	-

Table 38. Sample programs for Windows client systems

Description	Source	Executable module
CICS transaction	amqscic0.ccs	amqscicc
Header file for the CICS transaction	amqscih0.h	-
Tuxedo client program to put messages	amqstxpx.c	-
Tuxedo client program to get messages	amqstxgx.c	-
Tuxedo server program for the two client programs	amqstxsx.c	-
UBBCONFIG file for the Tuxedo programs	ubbstxcx.cfg	-
Field table file for the Tuxedo programs	amqstxvx.fld	-
View description file for the Tuxedo programs	amqstxvx.v	-
Makefile for the Tuxedo programs	amqstxmc.mak	-
ENVFILE file for the Tuxedo programs	amqstxen.env	-

Error message AMQ5203, as modified for CICS and Tuxedo applications:

UNIX

Windows

When you run CICS or Tuxedo applications that use an extended transactional client, you might see standard diagnostic messages. One of these has been modified for use with an extended transactional client

The messages that you might see in the IBM MQ error log files are documented in Diagnostic messages: AMQ4000-9999. Message AMQ5203 has been modified for use with an extended transactional client. Here is the text of the modified message:

AMQ5203: An error occurred calling the XA interface.

Explanation

The error number is &2 where a value of 1 indicates the supplied flags value of &1 was invalid, 2 indicates that there was an attempt to use threaded and non-threaded libraries in the same process, 3 indicates that there was an error with the supplied queue manager name '&3', 4 indicates that the resource manager ID of &1 was invalid, 5 indicates that an attempt was made to use a second queue manager called '&3' when another queue manager was already connected, 6 indicates that the Transaction Manager has been called when the application isn't connected to a queue manager, 7 indicates that the XA call was made while another call was in progress, 8 indicates that the xa_info string '&4' in the xa_open call contained an invalid parameter value for parameter name '&5', and 9 indicates that the xa_info string '&4' in the xa_open call is missing a required parameter, parameter name '&5'.

User response

Correct the error and try the operation again.

Preparing and running Microsoft Transaction Server applications

To prepare an MTS application to run as an IBM MQ MQI client application, follow these instructions as appropriate for your environment.

For general information about how to develop Microsoft Transaction Server (MTS) applications that access IBM MQ resources, see the section on MTS in the IBM MQ Help Center.

To prepare an MTS application to run as an IBM MQ MQI client application, do one of the following for each component of the application:

- If the component uses the C language bindings for the MQI, follow the instructions in “Preparing C programs in Windows” on page 414 but link the component with the library mqicxa.lib instead of mqic.lib.
- If the component uses the IBM MQ C++ classes, follow the instructions in “Building C++ programs on Windows” on page 820 but link the component with the library imqx23vn.lib instead of imqc23vn.lib.
- If the component uses the Visual Basic language bindings for the MQI, follow the instructions in the “Preparing Visual Basic programs in Windows” on page 417 but when you define the Visual Basic project, type MqType=3 in the **Conditional Compilation Arguments** field.
- If the component uses the IBM MQ Automation Classes for ActiveX (MQAX), define an environment variable, GMQ_MQ_LIB, with the value mqic32xa.dll.

You can define the environment variable from within your application, or you can define it so that its scope is system wide. However, defining it as system wide can cause any existing MQAX application, that does not define the environment variable from within the application, to behave incorrectly.

Preparing and running IBM MQ JMS applications

You can run IBM MQ JMS applications in client mode, with WebSphere Application Server as your transaction manager. You might see certain warning messages.

To prepare and run IBM MQ JMS applications in client mode, with WebSphere Application Server as your transaction manager, follow the instructions in “Using IBM MQ classes for JMS” on page 824.

When you run an IBM MQ JMS client application, you might see the following warning messages:

MQJE080

Insufficient license units - run setmqcap

MQJE081

File containing the license unit information is in the wrong format - run setmqcap

MQJE082

File containing the license unit information could not be found - run setmqcap

User exits, API exits, and IBM MQ installable services

This topic contains links to information about using and developing these programs.

For an introduction to how you can use user exits, API exits, and installable services to extend queue manager facilities, see Extending queue manager facilities.

For information on writing and compiling exits and installable services, see the subtopics.

Related information:

[Channel-exit programs for MQI channels](#)

[API exit reference](#)

[Installable services interface reference information](#)

 [Installable services interface reference information on IBM i](#)

Writing exits and installable services on UNIX, Linux and Windows



You can write and compile exits without linking to any IBM MQ libraries on UNIX, Linux and Windows.

About this task

This topic applies to UNIX, Linux, and Windows systems only. For details on writing exits and installable services for other platforms, see the relevant platform specific topics.

If IBM MQ is installed in a non-default location you must write and compile your exits without linking to any IBM MQ libraries.

You can write and compile exits on UNIX, Linux, and Windows systems without linking any of these IBM MQ libraries:

- mqmzf
- mqm
- mqmvx
- mqmvxd
- mqic
- mqutl

Existing exits that are linked to these libraries continue to work, providing that on UNIX and Linux systems IBM MQ is installed in the default location.

Procedure

1. Include the `cmqec.h` header file. Including this header file automatically includes the `cmqc.h`, `cmqxc.h` and `cmqzc.h` header files.
2. Write the exit so that MQI and DCI calls are made through the MQIEP structure. For more information about the MQIEP structure, see MQIEP structure.
 - Installable services
 - Use the **Hconfig** parameter to point to the MQZEP call.
 - You must check that the first 4 bytes of **Hconfig** match the **StrucId** of the MQIEP structure before using the **Hconfig** parameter.
 - For more information about writing installable service components, see MQIEP.
 - API exits
 - Use the **Hconfig** parameter to point to the MQXEP call.
 - You must check that the first 4 bytes of **Hconfig** match the **StrucId** of the MQIEP structure before using the **Hconfig** parameter.
 - For more information about writing API exits, see “Writing API exits” on page 333.
 - Channel exits
 - Use the **pEntryPoints** parameter of the MQCXP structure to point to MQI and DCI calls.
 - You must check that the MQCXP version number is at version 8 or higher before using **pEntryPoints**.
 - For more information about writing channel exits, see “Writing channel-exit programs” on page 344.
 - Data conversion exits
 - Use the **pEntryPoints** parameter of the MQDXP structure to point to MQI and DCI calls.
 - You must check that the MQDXP version number is at version 2 or higher before using **pEntryPoints**.
 - You can use the **crtmqcvx** command and the `amqsvfc0.c` source file to create data conversion code that uses the **pEntryPoints** parameter. See “Writing a data-conversion exit for IBM MQ for Windows” on page 372 and “Writing a data-conversion exit for IBM MQ on UNIX and Linux systems” on page 369.
 - If you have existing data conversion exits that were generated using the **crtmqcvx** command, you must regenerate the exit using the updated command.
 - For more information about writing data conversion exits, see “Writing data-conversion exits” on page 364.
 - Pre-connect exits
 - Use the **pEntryPoints** parameter of the MQNXP structure to point to MQI and DCI calls.
 - You must check that the MQNXP version number is at version 2 or higher before using **pEntryPoints**.
 - For more information about writing pre-connect exits, see “Referencing connection definitions using a pre-connect exit from a repository” on page 374.
 - Publish exits
 - Use the **pEntryPoints** parameter of the MQPSXP structure to point to MQI and DCI calls.
 - You must check that the MQPSXP version number is at version 2 or higher before using **pEntryPoints**.
 - For more information about writing publish exits, see “Writing and compiling publish exits” on page 377.
 - Cluster workload exits

- Use the **pEntryPoints** parameter of the MQWXP structure to point to MQXCLWLN calls.
- You must check that the MQWXP version number is at version 4 or higher before using **pEntryPoints**.
- For more information about writing cluster workload exits, see “Writing and compiling cluster workload exits” on page 379.

For example, in a channel exit calling MQPUT:

```
pChannelExitParms -> pEntryPoints -> MQPUT_Call(pChannelExitParms -> Hconn,
                                                Hobj,
                                                &md,
                                                &pmo,
                                                messLen,
                                                buffer,
                                                &CompCode,
                                                &Reason);
```

Further examples can be seen in the “Using the IBM MQ sample procedural programs” on page 457.

3. Compile the exit:

- Do not link to the IBM MQ libraries.
- Do not include an embedded RPath to any IBM MQ libraries in your exit.
- For more information about compiling your exit, see one of the following topics:
 - API exits: “Compiling API exits” on page 335.
 - Channel exits, publish exits, Cluster workload exits: “Compiling channel exit programs on Windows, UNIX and Linux systems” on page 363.
 - Data conversion exits: “Writing data-conversion exits” on page 364.

4. Put the exit in one of the following places:

- A path of your choosing that you fully qualify when configuring the exit
- The default exit path, in a specific installation directory. For example, *MQ_DATA_PATH/exits/installation2*.
- The default exit path

The default exit path is *MQ_DATA_PATH/exits* for 32 bit exits, and *MQ_DATA_PATH/exits64* for 64 bit exits. You can change these paths in the *qm.ini* or *mqlclient.ini* file. For more information, see Exit path. On Windows and Linux, you can use the IBM MQ Explorer to change the path:

- a. Right click the queue manager name
- b. Click **Properties...**
- c. Click **Exits**
- d. In the exits default path field, specify the path name of the directory that holds the exit program.

If an exit is placed in both a specific installation directory and in the default path directory, the specific installation directory exit is used by the installation of IBM MQ named in the path. For example, the exit is placed in */exits/installation2* and in */exits*, but not in */exits/installation1*. The IBM MQ installation *installation2* uses the exit from */exits/installation2*. The IBM MQ installation *installation1* uses the exit from the */exits* directory.

5. If necessary, configure the exit:

- Installable services: “Configuring services and components” on page 322.
- API exits: “Configuring API exits” on page 339.
- Channel exits: “Configuring channel exits” on page 364.
- Publish exits: “Configuring publish exits” on page 378.
- Pre-connect exits: PreConnect stanza of the client configuration file.

API exits not linked with an MQI library:

Under certain circumstances, you should link your existing API exit, that cannot be re-coded to use the MQIEP function pointers, with an IBM MQ API library.

This is necessary, so that your existing API exit can be successfully loaded, by the runtime linker of your system, into programs that do not already have the function pointers loaded.

Note: This information is limited to those existing API exits that make MQI calls directly. That is, those exits that do not use , MQIEP. Where possible, you should plan to re-code the exit to use the MQIEP entry points instead.

From IBM MQ Version 8.0, **runmqsc** is an example of a program that does not link directly with an MQI library.

Therefore, an API exit that has not been linked with its required IBM MQ API library, or re-coded to use the MQIEP, fails to load into **runmqsc**.

You see errors in the queue manager error log, for example, AMQ6175: The system could not dynamically load the shared library, together with qualifying text such as undefined symbol: MQCONN.

and AMQ7214: The module for API Exit 'myexitname' could not be loaded.

Related tasks:

“Writing exits and installable services on UNIX, Linux and Windows” on page 311

You can write and compile exits without linking to any IBM MQ libraries on UNIX, Linux and Windows.

Installable services and components for UNIX, Linux and Windows



This section introduces the installable services and the functions and components associated with them. The interface to these functions is documented so that you, or software vendors, can supply components.

The main reasons for providing IBM MQ installable services are:

- To provide you with the flexibility of choosing whether to use components provided by IBM MQ products, or replace or augment them with others.
- To allow vendors to participate, by providing components that might use new technologies, without making internal changes to IBM MQ products.
- To allow IBM MQ to exploit new technologies faster and cheaper, and so provide products earlier and at lower prices.

Installable services and *service components* are part of the IBM MQ product structure. At the center of this structure is the part of the queue manager that implements the function and rules associated with the Message Queue Interface (MQI). This central part requires a number of service functions, called *installable services*, in order to perform its work. The installable services are:

- Authorization service
- Name service

Each installable service is a related set of functions implemented using one or more *service components*. Each component is invoked using a properly-architected, publicly-available interface. This enables independent software vendors and other third parties to provide installable components to augment or replace those provided by the IBM MQ products. Table 39 on page 315 summarizes the services and components that can be used.

Table 39. Installable service components summary

Installable service	Supplied component	Function	Requirements
Authorization service	object authority manager (OAM)	Provides authorization checking on commands and MQI calls. Users can write their own component to augment or replace the OAM. For example, to check that a user ID has authority to open a queue.	(Appropriate platform authorization facilities are assumed)
Name service	None	Provides support to the queue manager for looking up the name of the queue manager that owns a specified queue. <ul style="list-style-type: none"> User defined Note: Shared queues must have their Scope attribute set to CELL.	<ul style="list-style-type: none"> A third-party or user-written name manager

The installable services interface is described in Installable services interface reference information.

Related information:

Configuring installable services

Writing a service component:

This section describes the relationship between services, components, entry points, and return codes.

Functions and components

Each service consists of a set of related functions. For example, the name service contains function for:

- Looking up a queue name and returning the name of the queue manager where the queue is defined
- Inserting a queue name into the service's directory
- Deleting a queue name from the service's directory

It also contains initialization and termination functions.

An installable service is provided by one or more service components. Each component can perform some or all of the functions that are defined for that service. For example, in IBM MQ for AIX, the supplied authorization service component, the OAM, performs all the available functions. See "Authorization service interface" on page 319 for more information. The component is also responsible for managing any underlying resources or software (for example, an LDAP directory) that it needs to implement the service. Configuration files provide a standard way of loading the component and determining the addresses of the functional routines that it provides.

Figure 52 on page 316 shows how services and components are related:

- A service is defined to a queue manager by stanzas in a configuration file.
- Each service is supported by supplied code in the queue manager. Users cannot change this code and therefore cannot create their own services.
- Each service is implemented by one or more components; these can be supplied with the product or user-written. Multiple components for a service can be invoked, each supporting different facilities within the service.
- Entry points connect the service components to the supporting code in the queue manager.

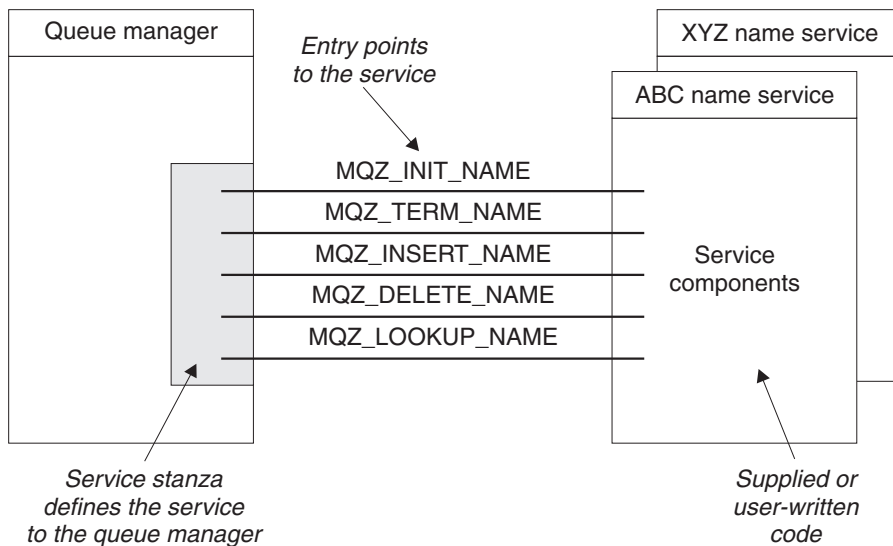


Figure 52. Understanding services, components, and entry points

Entry-points

Each service component is represented by a list of the entry-point addresses of the routines that support a particular installable service. The installable service defines the function to be performed by each routine.

The ordering of the service components when they are configured defines the order in which entry-points are called in an attempt to satisfy a request for the service.

In the supplied header file `cmqzc.h`, the supplied entry points to each service have an `MQZID_` prefix.

If the services are present, the services are loaded in a predefined order. The following list shows the services, and the order in which they are initialized.

1. NameService
2. AuthorizationService
3. UserIdentifierService

The AuthorizationService is the only service that is configured by default. Configure the NameService and UserIdentifierService manually if you want to use them.

Services, and service components have a one-to-one or one-to-many mapping. Multiple service components can be defined for each service. On UNIX and Linux systems, the ServiceComponent stanza's Service value must match the Service stanza's Name value in the `qm.ini` file. On Windows, the ServiceComponent's Service registry key value must match the Name registry key value, and is defined as: `HKEY_LOCAL_MACHINE\SOFTWARE\IBM\WebSphere MQ\Installation\MQ_INSTALLATION_NAME\Configuration\QueueManager\qmname` where `qmname` is the name of the queue manager.

For UNIX and Linux systems, service components are started in the order they are defined in the `qm.ini` file. On Windows, because the Windows registry is used, IBM MQ issues a `RegEnumKey` call which returns the values in alphabetic order. Therefore, on Windows the services are called in alphabetic order, as they are defined in the registry.

The ordering of the ServiceComponent definitions is significant. This ordering dictates the order in which components are run for a given service. For example, the AuthorizationService on Windows is configured with the default OAM component named `MQSeries.WindowsNT.auth.service`. Additional

components can be defined for this service in order to override the default OAM. Unless `MQCACF_SERVICE_COMPONENT` is specified, the first component encountered in alphabetic order is used to process the request, and the name for that component is used.

Return codes

Service components provide return codes to the queue manager to report on various conditions. They report the success or failure of the operation, and indicate whether the queue manager is to proceed to the next service component. A separate *Continuation* parameter carries this indication.

Component data

A single service component might require data to be shared between its various functions. Installable services provide an optional data area to be passed on each invocation of a service component. This data area is for the exclusive use of the service component. It is shared by all the invocations of a particular function, even if they are made from different address spaces or processes. It is guaranteed to be addressable from the service component whenever it is called. You must declare the size of this area in the *ServiceComponent* stanza.

Initialization and termination of components:

The use of component initialization and termination options.

When the component initialization routine is invoked, it must call the queue manager **MQZEP** function for each entry-point supported by the component. **MQZEP** defines an entry-point to the service. All the undefined exit points are assumed to be NULL.

A component is always invoked once with the primary initialization option, before it is invoked in any other way.

A component can be invoked with the secondary initialization option on certain platforms. For example, it can be invoked once for each operating system process, thread, or task by which the service is accessed.

If secondary initialization is used:

- The component can be invoked more than once for secondary initialization. For each such call, a matching call for secondary termination is issued when the service is no longer needed.
For naming services this is the `MQZ_TERM_NAME` call.
For authorization services this is the `MQZ_TERM_AUTHORITY` call.
- The entry points must be re-specified (by calling **MQZEP**) each time the component is called for primary and secondary initialization.
- Only one copy of component data is used for the component; there is not a different copy for each secondary initialization.
- The component is not invoked for any other calls to the service (from the operating system process, thread, or task, as appropriate) before secondary initialization has been carried out.
- The component must set the **Version** parameter to the same value for primary and secondary initialization.

The component is always invoked with the primary termination option once, when it is no longer required. No further calls are made to this component.

The component is invoked with the secondary termination option, if it has been invoked for secondary initialization.

Object authority manager (OAM):

The authorization service component supplied with the IBM MQ products is called the Object Authority Manager (OAM).

By default, the OAM is active and works with the control commands **dspmqa** (display authority), **dmpmqa** (dump authority), and **setmqa** (set or reset authority).

The syntax of these commands and how to use them are described in *The control commands*.

The OAM works with the *entity* of a principal or group:

- **UNIX** **Linux** On UNIX and Linux systems, a principal is a user ID, or an ID associated with an application program running on behalf of a user; a group is a system-defined collection of principals.
- **Windows** On Windows systems, a principal is a Windows user ID, or an ID associated with an application program running on behalf of a user; a group is a Windows group.

Authorizations can be granted or revoked at the principal or group level.

When an MQI request is made or a command is issued, the OAM checks whether the entity associated with the operation has authorization to perform the requested operation and to access the specified queue manager resources.

The authorization service enables you to augment or replace the authority checking provided for queue managers by writing your own authorization service component.

Name service:

The name service is an installable service that provides support to the queue manager for looking up the name of the queue manager that owns a specified queue. No other queue attributes can be retrieved from a name service.

The name service enables an application to open remote queues for output as if they were local queues. A name service is not invoked for objects other than queues.

Note: The remote queues must have their **Scope** attribute set to CELL.

When an application opens a queue, it looks for the name of the queue first in the queue manager's directory. If it does not find it there, it looks in as many name services as have been configured, until it finds one that recognizes the queue name. If none recognizes the name, the open fails.

The name service returns the owning queue manager for that queue. The queue manager then continues with the MQOPEN request as if the command had specified the queue and queue manager name in the original request.

The name service interface (NSI) is part of the IBM MQ framework.

How the name service works

If a queue definition specifies the **Scope** attribute as queue manager, that is, SCOPE(QMGR) in MQSC, the queue definition (along with all the queue attributes) is stored in the queue manager's directory only. This cannot be replaced by an installable service.

If a queue definition specifies the **Scope** attribute as cell, that is, SCOPE(CELL) in MQSC, the queue definition is again stored in the queue manager's directory, along with all the queue attributes. However,

the queue and queue manager name are also stored in a name service. If no service is available that can store this information, a queue with the *Scope* cell cannot be defined.

The directory in which the information is stored can be managed by the service, or the service can use an underlying service, for example, an LDAP directory, for this purpose. In either case, definitions stored in the directory must persist, even after the component and queue manager have terminated, until they are explicitly deleted.

Note:

1. To send a message to a remote host's local queue definition (with a scope of CELL) on a different queue manager within a naming directory cell, you need to define a channel.
2. You cannot get messages directly from the remote queue, even when it has a scope of CELL.
3. No remote queue definition is required when sending to a queue with a scope of CELL.
4. The naming service centrally defines the destination queue, although you still need a transmission queue to the destination queue manager and a pair of channel definitions. In addition, the transmission queue on the local system must have the same name as the queue manager owning the target queue, with the scope of cell, on the remote system.

For example, if the remote queue manager has the name QM01, the transmission queue on the local system must also have the name QM01.

Authorization service interface:

The authorization service provides entry points for use by the queue manager.

The entry points are as follows:

MQZ_AUTHENTICATE_USER

Authenticates a user ID and password, and can set identity context fields.

MQZ_CHECK_AUTHORITY

Checks whether an entity has authority to perform one or more operations on a specified object.

MQZ_CHECK_PRIVILEGED

Checks whether a specified user is a privileged user.

MQZ_COPY_ALL_AUTHORITY

Copies all the current authorizations that exist for a referenced object to another object.

MQZ_DELETE_AUTHORITY

Deletes all authorizations associated with a specified object.

MQZ_ENUMERATE_AUTHORITY_DATA

Retrieves all the authority data that matches the selection criteria specified.

MQZ_FREE_USER

Frees associated allocated resources.

MQZ_GET_AUTHORITY

Gets the authority that an entity has to access a specified object.

MQZ_GET_EXPLICIT_AUTHORITY

Gets either the authority that a named group has to access a specified object (but without the additional authority of the **nobody** group) or the authority that the primary group of the named principal has to access a specified object.

MQZ_INIT_AUTHORITY

Initializes authorization service component.

MQZ_INQUIRE

Queries the supported functionality of the authorization service.

MQZ_REFRESH_CACHE

Refresh all authorizations.

MQZ_SET_AUTHORITY

Sets the authority that an entity has to a specified object.

MQZ_TERM_AUTHORITY

Terminates authorization service component.

In addition, on IBM MQ for Windows, the authorization service provides the following entry points for use by the queue manager:

- **MQZ_CHECK_AUTHORITY_2**
- **MQZ_GET_AUTHORITY_2**
- **MQZ_GET_EXPLICIT_AUTHORITY_2**
- **MQZ_SET_AUTHORITY_2**

These entry points support the use of the Windows Security Identifier (NT SID).

These names are defined as **typedef s**, in the header file `cmqzc.h`, which can be used to prototype the component functions.

The initialization function (**MQZ_INIT_AUTHORITY**) must be the main entry point for the component. The other functions are invoked through the entry point address that the initialization function has added into the component entry point vector.

Name service interface:

A name service provides entry points for use by the queue manager.

The following entry points are provided:

MQZ_INIT_NAME

Initialize the name service component.

MQZ_TERM_NAME

Terminate the name service component.

MQZ_LOOKUP_NAME

Look up the queue manager name for the specified queue.

MQZ_INSERT_NAME

Insert an entry containing the owning queue manager name for the specified queue into the directory used by the service.

MQZ_DELETE_NAME

Delete the entry for the specified queue from the directory used by the service.

If there is more than one name service configured:

- For lookup, the **MQZ_LOOKUP_NAME** function is invoked for each service in the list until the queue name is resolved (unless any component indicates that the search should stop).
- For insert, the **MQZ_INSERT_NAME** function is invoked for the first service in the list that supports this function.
- For delete, the **MQZ_DELETE_NAME** function is invoked for the first service in the list that supports this function.

Do not have more than one component that supports the insert and delete functions. However, a component that only supports lookup is feasible, and could be used, for example, as the last component in the list to resolve any name that is not known by any other name service component to a queue manager at which the name can be defined.

In the C programming language the names are defined as function data types using the typedef statement. These can be used to prototype the service functions, to ensure that the parameters are correct.

The header file that contains all the material specific to installable services is `cmqzc.h` for the C language.

Apart from the initialization function (`MQZ_INIT_NAME`), which must be the component's main entry point, functions are invoked by the entry point address that the initialization function has added, using the `MQZEP` call.

Using multiple service components:

You can install more than one component for a service. This allows components to provide only partial implementations of the service, and to rely on other components to provide the remaining functions.

Example of using multiple components

Suppose you create two a name services components called `ABC_name_serv` and `XYZ_name_serv`.

ABC_name_serv

This component supports inserting a name in, or deleting a name from, the service directory, but does not support looking up a queue name.

XYZ_name_serv

This component supports looking up a queue name, but does not support inserting a name in, or deleting a name from, the service directory.

Component `ABC_name_serv` holds a database of queue names, and uses two simple algorithms to either insert, or delete, a name from the service directory.

Component `XYZ_name_serv` uses a simple algorithm that returns a fixed queue manager name for any queue name with which it is invoked. It does not hold a database of queue names, and therefore does not support the insert and delete functions.

The components are installed on the same queue manager. The *ServiceComponent* stanzas are ordered so that component `ABC_name_serv` is invoked first. Any calls to insert or delete a queue in a component directory are handled by component `ABC_name_serv` ; it is the only one that implements these functions. However, a lookup call that component `ABC_name_serv` cannot resolve is passed on to the lookup-only component, `XYZ_name_serv`. This component supplies a queue manager name from its simple algorithm.

Omitting entry points when using multiple components

If you decide to use multiple components to provide a service, you can design a service component that does not implement certain functions. The installable services framework places no restrictions on which you can omit. However, for specific installable services, omission of one or more functions might be logically inconsistent with the purpose of the service.

Example of entry points used with multiple components

Table 40 on page 322 shows an example of the installable name service for which the two components have been installed. Each supports a different set of functions associated with this particular installable service. For insert function, the ABC component entry-point is invoked first. Entry points that have not

been defined to the service (using **MQZEP**) are assumed to be NULL. An entry-point for initialization is provided in the table, but this is not required because initialization is carried out by the main entry-point of the component.

When the queue manager has to use an installable service, it uses the entry-points defined for that service (the columns in Table 40). Taking each component in turn, the queue manager determines the address of the routine that implements the required function. It then calls the routine, if it exists. If the operation is successful, any results and status information are used by the queue manager.

Table 40. Example of entry-points for an installable service

Function number	ABC name service component	XYZ name service component
MQZID_INIT_NAME (Initialize)	ABC_initialize()	XYZ_initialize()
MQZID_TERM_NAME (Terminate)	ABC_terminate()	XYZ_terminate()
MQZID_INSERT_NAME (Insert)	ABC_Insert()	NULL
MQZID_DELETE_NAME (Delete)	ABC_Delete()	NULL
MQZID_LOOKUP_NAME (Lookup)	NULL	XYZ_Lookup()

If the routine does not exist, the queue manager repeats this process for the next component in the list. In addition, if the routine does exist but returns a code indicating that it could not perform the operation, the attempt continues with the next available component. Routines in service components might return a code that indicates that no further attempts to perform the operation should be made.

Configuring services and components:

Configure service components using the queue manager configuration files, except on Windows systems, where each queue manager has its own stanza in the Registry.

1. Add stanzas to the queue manager configuration file to define the service to the queue manager and specify the location of the module.

Each service used must have a *Service* stanza, which defines the service to the queue manager.

For each component within a service, there must be a *ServiceComponent* stanza. This identifies the name and path of the module containing the code for that component.

For more information, see “Service stanza format” on page 323 and “Service component stanza format” on page 323

The authorization service component, known as the Object Authority Manager (OAM), is supplied with the product. When you create a queue manager, the queue manager configuration file (or the Registry on Windows systems) is automatically updated to include the appropriate stanzas for the authorization service and for the default component (the OAM). For the other components, you must configure the queue manager configuration file manually.

The code for each service component is loaded into the queue manager when the queue manager is started, using dynamic binding, where this is supported on the platform.

2. Stop and restart the queue manager to activate the component.

Service stanza format:

The Service stanza contains the name of the service and the number of entry points defined for the service.

The format of the stanza is as follows:

```
Service:
  Name=service_name
  EntryPoints=entries
  SecurityPolicy=policy
```

where:



service_name


The name of the service. This is defined by the service.

entries

The number of entry points defined for the service. This includes the initialization and termination entry points.

policy

  On UNIX and Linux systems: user, group, or default. 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, the default value used is group-based authorization. Restart the queue manager for changes to become effective. See also “Configuring authorization service stanzas on UNIX and Linux” on page 324.

 On Windows systems: NTSDsRequired (the Windows Security Identifier), or Default. If you do not specify NTSDsRequired, the Default value is used. This attribute is valid only if **Name** has a value of AuthorizationService. See also “Configuring authorization service stanzas on Windows” on page 325.

Service component stanza format:

The **Service** and **ServiceComponent** stanzas can occur in any order.

The format of the service component stanza is:

```
ServiceComponent:
  Service=service_name
  Name=component_name
  Module=module_name
  ComponentDataSize=size
```

where:

service_name

The name of the service. This must match the Name specified in a service stanza.

component_name

A descriptive name of the service component. This must be unique, and contain only the 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 you use a name starting with a company trademark or similar distinguishing string.

module_name

The name of the module to contain the code for this component.

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.

The **Service** and **ServiceComponent** stanzas can occur in any order and the stanza keys under them can also occur in any order. For either of these stanzas, all the stanza keys must be present. If a stanza key is duplicated, the last one is used.

At startup time, the queue manager processes each service component entry in the configuration file in turn. It then loads the specified component module, invoking the entry point of the component (which must be the entry point for initialization of the component), passing it a configuration handle.

Configuring authorization service stanzas on UNIX and Linux: 

On UNIX and Linux, each queue manager has its own queue manager configuration file.

For example, the default path and file name of the queue manager configuration file for queue manager QMNAME is `/var/mqm/qmgrs/QMNAME/qm.ini`.

The *Service* stanza and the *ServiceComponent* stanza for the default authorization component are added to `qm.ini` automatically, but can be overridden by `mqsnoaut`. Any other *ServiceComponent* stanzas must be added manually.

For example, the following stanzas in the queue manager configuration file define two authorization service components on IBM MQ for AIX. `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.


```
Service:
  Name=AuthorizationService
  EntryPoints=13

ServiceComponent:
  Service=AuthorizationService
  Name=MQSeries.UNIX.auth.service
Module= MQ_INSTALLATION_PATH/lib/amqzfu
  ComponentDataSize=0

ServiceComponent:
  Service=AuthorizationService
  Name=user.defined.authorization.service
  Module=/usr/bin/udas01
  ComponentDataSize=96
```

Figure 53. UNIX and Linux authorization service stanzas in `qm.ini`

The service component stanza (`MQSeries.UNIX.auth.service`) defines the default authorization service component, the OAM. If you remove this stanza and restart the queue manager, the OAM is disabled and no authorization checks are made.

Configuring authorization service stanzas on Windows: 

On IBM MQ for Windows each queue manager has its own stanza in the registry.

The *Service* stanza and the *ServiceComponent* stanza for the default authorization component are added to the Registry automatically, but can be overridden using `mqsnout`. Any other *ServiceComponent* stanzas must be added manually.

You can also add the `SecurityPolicy` attribute using the IBM MQ services. The `SecurityPolicy` attribute applies only if the service specified on the *Service* stanza is the authorization service, that is, the default OAM. The `SecurityPolicy` attribute allows you to specify the security policy for each queue manager. The possible values are:

Default



Specify `Default` if you want 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

Requires that an NT SID is passed to the OAM when performing security checks.

For information about the *Service* stanza format, see “Service stanza format” on page 323. For more general information about security, see *Setting up security on Windows, UNIX and Linux systems*.

The service component stanza, `MQSeries.WindowsNT.auth.service` defines the default authorization service component, the OAM. If you remove this stanza and restart the queue manager, the OAM is disabled and no authorization checks are made.

Configuring name service stanzas: UNIX and Linux systems:  


On UNIX and Linux systems, each queue manager has its own queue manager configuration file.

The following examples of UNIX and Linux configuration file stanzas for the name service specify a name service component provided by the (fictitious) ABC company.

```
# Stanza for name service
Service:
  Name=NameService
  EntryPoints=5

# Stanza for name service component, provided by ABC
ServiceComponent:
  Service=NameService
  Name=ABC.Name.Service
  Module=/usr/lib/abcname
  ComponentDataSize=1024
```

Figure 54. Name service stanzas in `qm.ini` (for UNIX and Linux systems)

Note:  On Windows systems, name service stanza information is stored in the Registry.

Refreshing the OAM after changing a user's authorization:

In IBM MQ, you can refresh the OAM's authorization group information immediately after changing a user's authorization group membership, reflecting changes made at the operating system level, without needing to stop and restart the queue manager. In order to do this, issue the **REFRESH SECURITY** command.

Note: When you change authorizations with the `setmqaut` command, the OAM implements such changes immediately.

Queue managers store authorization data on a local queue called `SYSTEM.AUTH.DATA.QUEUE`. This data is managed by `amqzfuma.exe`.

Related information:

REFRESH SECURITY

Installable services and components on IBM i



Use this information to learn about the installable services and the functions and components associated with them. The interface to these functions is documented so that you, or software vendors, can supply components.

The main reasons for providing IBM MQ installable services are:

- To provide you with the flexibility of choosing whether to use components provided by IBM MQ for IBM i, or replace or augment them with others.
- To allow vendors to participate, by providing components that might use new technologies, without making internal changes to IBM MQ for IBM i.
- To allow IBM MQ to exploit new technologies faster and cheaper, and so provide products earlier and at lower prices.

Installable services and *service components* are part of the IBM MQ product structure. At the center of this structure is the part of the queue manager that implements the function and rules associated with the Message Queue Interface (MQI). This central part requires a number of service functions, called *installable services*, in order to perform its work. The installable service available in IBM MQ for IBM i is the authorization service.

Each installable service is a related set of functions implemented using one or more *service components*. Each component is invoked using a properly-architected, publicly-available interface. This enables independent software vendors and other third parties to provide installable components to augment or replace those provided by IBM MQ for IBM i. Table 41 summarizes support for the authorization service.

Table 41. Authorization service components summary

Supplied component	Function	Requirements
Object authority manager (OAM)	Provides authorization checking on commands and MQI calls. Users can write their own component to augment or replace the OAM.	(Appropriate platform authorization facilities are assumed)
DCE name service component Note: DCE is only supported on versions of IBM MQ earlier than V6.0.	<ul style="list-style-type: none"> • Allows queue managers to share queues, or • User defined Note: Shared queues must have their Scope attribute set to CELL .	<ul style="list-style-type: none"> • DCE is required for the supplied component, or • A third-party or user-written name manager

Functions and components on IBM i:



Use this information to understand the functions and components, entry-points, return codes, and component data that you can use in IBM MQ for IBM i.

Each service consists of a set of related functions. For example, the name service contains function for:

- Looking up a queue name and returning the name of the queue manager where the queue is defined
- Inserting a queue name into the service's directory
- Deleting a queue name from the service's directory

It also contains initialization and termination functions.

An installable service is provided by one or more service components. Each component can perform some or all of the functions that are defined for that service. The component is also responsible for managing any underlying resources or software that it needs to implement the service. Configuration files provide a standard way of loading the component and determining the addresses of the functional routines that it provides.

Services and components are related as follows:

- A service is defined to a queue manager by stanzas in a configuration file.
- Each service is supported by supplied code in the queue manager. Users cannot change this code and therefore cannot create their own services.
- Each service is implemented by one or more components; these can be supplied with the product or user-written. Multiple components for a service can be invoked, each supporting different facilities within the service.
- Entry points connect the service components to the supporting code in the queue manager.

Entry-points

Each service component is represented by a list of the entry-point addresses of the routines that support a particular installable service. The installable service defines the function to be performed by each routine. The ordering of the service components when they are configured defines the order in which entry-points are called in an attempt to satisfy a request for the service. In the supplied header file `cmqzc.h`, the supplied entry points to each service have an `MQZID_` prefix.

Return codes

Service components provide return codes to the queue manager to report on a variety of conditions. They report the success or failure of the operation, and indicate whether the queue manager is to proceed to the next service component. A separate *Continuation* parameter carries this indication.

Component data

A single service component might require data to be shared between its various functions. Installable services provide an optional data area to be passed on each invocation of a particular service component. This data area is for the exclusive use of the service component. It is shared by all the invocations of a given function, even if they are made from different address spaces or processes. It is guaranteed to be addressable from the service component whenever it is called. You must declare the size of this area in the *ServiceComponent* stanza.

Initialization on IBM i:

When the component initialization routine is invoked, it must call the queue manager MQZEP function for each entry-point supported by the component. MQZEP defines an entry-point to the service. All the undefined exit points are assumed to be NULL.

Primary initialization

A component is always invoked with this option once, before it is invoked in any other way.

Secondary initialization

A component can be invoked with this option on certain platforms. For example, it can be invoked once for each operating system process, thread, or task by which the service is accessed.

If secondary initialization is used:

- The component can be invoked more than once for secondary initialization. For each such call, a matching call for secondary termination is issued when the service is no longer needed. For authorization services this is the MQZ_TERM_AUTHORITY call.
- The entry points must be re-specified (by calling MQZEP) each time the component is called for primary and secondary initialization.
- Only one copy of component data is used for the component; there is not a different copy for each secondary initialization.
- The component is not invoked for any other calls to the service (from the operating system process, thread, or task, as appropriate) before secondary initialization has been carried out.
- The component must set the **Version** parameter to the same value for primary and secondary initialization.

Primary termination

The component is always started with this option once, when it is no longer required. No further calls are made to this component.

Secondary termination

The component is started with this option, if it has been started for secondary initialization.

Configuring services and components on IBM i:

Configure service components using the queue manager configuration files. Each service used must have a *Service* stanza, which defines the service to the queue manager.

For each component within a service, there must be a *ServiceComponent* stanza. This identifies the name and path of the module containing the code for that component.

The authorization service component, known as the object authority manager (OAM), is supplied with the product. When you create a queue manager, the queue manager configuration file is automatically updated to include the appropriate stanzas for the authorization service and for the default component (the OAM).

The code for each service component is loaded into the queue manager when the queue manager is started, using dynamic binding, where this is supported on the platform.

Service stanza format

The format of the **Service** stanza is:

```
Service:  
  Name=service_name  
  EntryPoints=entries
```


where:

service_name

The name of the service. This is defined by the service.

entries

The number of entry points defined for the service. This includes the initialization and termination entry points.

Service component stanza format

The format of the **Service component** stanza is:

```
ServiceComponent:  
  Service=service_name  
  Name=component_name  
  Module=module_name  
  ComponentDataSize=size
```

where:

service_name

The name of the service. This must match the *Name* specified in a service stanza.

component_name

A descriptive name of the service component. This must be unique, and contain only the 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 you use a name starting with a company trademark or similar distinguishing string.

module_name

The name of the module to contain the code for this component. Specify a full path name.

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.

These two stanzas can occur in any order and the stanza keys under them can also occur in any order. For either of these stanzas, all the stanza keys must be present. If a stanza key is duplicated, the last one is used.

At startup time, the queue manager processes each service component entry in the configuration file in turn. It then loads the specified component module, invoking the entry-point of the component (which must be the entry-point for initialization of the component), passing it a configuration handle.

Creating your own service component on IBM i:

Use this information to learn how to create a service component on IBM MQ for IBM i.

To create your own service component:

- Ensure that the header file `cmqzc.h` is included in your program.
- Create the shared library by compiling the program and linking it with the shared libraries `libmqm*` and `libmqmzf*`.

Note: Because the agent can run in a threaded environment, you must build the OAM to run in a threaded environment. This includes using the threaded versions of `libmqm` and `libmqmzf`.

- Add stanzas to the queue manager configuration file to define the service to the queue manager and to specify the location of the module.
- Stop and restart the queue manager to activate the component.

Authorization service on IBM i:

The authorization service is an installable service that enables queue managers to invoke authorization facilities, for example, checking that a user ID has authority to open a queue.

This service is a component of the IBM MQ security enabling interface (SEI), which is part of the IBM MQ framework. The following subjects are discussed:

- “Object authority manager (OAM)”
- “Defining the service to the operating system”
- “Configuring authorization service stanzas”
- “Authorization service interface on IBM i” on page 331

Object authority manager (OAM)

The authorization service component supplied with the IBM MQ products is called the object authority manager (OAM). By default, the OAM is active and works with the following control commands:

- **WRKMQAUT** work with authority
- **WRKMQAUTD** work with authority data
- **DSPMQAUT** display object authority
- **GRTMQAUT** grant object authority
- **RVKMQAUT** revoke object authority
- **RFRMQAUT** refresh security

The syntax of these commands and how to use them are described in the CL command help. The OAM works with the *entity* of a principal or group.

When an MQI request is made or a command is issued, the OAM checks the authorization of the entity associated with the operation to see whether it can do the following actions:

- Perform the requested operation.
- Access the specified queue manager resources.

The authorization service enables you to augment or replace the authority checking provided for queue managers by writing your own authorization service component.

Defining the service to the operating system

The authorization service stanzas in the queue manager configuration file `qm.ini` define the authorization service to the queue manager. See “Configuring services and components on IBM i” on page 328 for information about the types of stanza.

Configuring authorization service stanzas

On IBM MQ for IBM i:

Principal

Is an IBM i system user profile.

Group Is an IBM i system group profile.

Authorizations can be granted or revoked at the group level only. A request to grant or revoke a user's authority updates the primary group for that user.

Each queue manager has its own queue manager configuration file. For example, the default path and file name of the queue manager configuration file for queue manager `QMNAME` is `/QIBM/UserData/mqm/qmgrs/QMNAME/qm.ini`.

The *Service* stanza and the *ServiceComponent* stanza for the default authorization component are added to `qm.ini` automatically, but can be overridden by `WRKENVVAR`. Any other *ServiceComponent* stanzas must be added manually.

For example, the following stanzas in the queue manager configuration file define two authorization service components:

```
Service:
  Name=AuthorizationService
  EntryPoints=7

ServiceComponent:
  Service=AuthorizationService
  Name=MQ.UNIX.authorization.service
  Module=QMQM/AMQZFU
  ComponentDataSize=0

ServiceComponent:
  Service=AuthorizationService
  Name=user.defined.authorization.service
  Module=LIBRARY/SERVICE PROGRAM NAME
  ComponentDataSize=96
```

Figure 55. Authorization service stanzas in `qm.ini` on IBM i

The first service component stanza `MQ.UNIX.authorization.service` defines the default authorization service component, the OAM. If you remove this stanza and restart the queue manager, the OAM is disabled and no authorization checks are made.

Authorization service interface on IBM i: 

The authorization service interface provides several entry points for use by the queue manager.

MQZ_AUTHENTICATE_USER

Authenticates a user ID and password, and can set identity context fields.

MQZ_CHECK_AUTHORITY

Checks whether an entity has authority to perform one or more operations on a specified object.

MQZ_COPY_ALL_AUTHORITY

Copies all the current authorizations that exist for a referenced object to another object.

MQZ_DELETE_AUTHORITY

Deletes all authorizations associated with a specified object.

MQZ_ENUMERATE_AUTHORITY_DATA

Retrieves all the authority data that matches the selection criteria specified.

MQZ_FREE_USER

Frees associated allocated resources.

MQZ_GET_AUTHORITY

Gets the authority that an entity has to access a specified object.

MQZ_GET_EXPLICIT_AUTHORITY

Gets either the authority that a named group has to access a specified object (but without the additional authority of the **nobody** group) or the authority that the primary group of the named principal has to access a specified object.

MQZ_INIT_AUTHORITY

Initializes authorization service component.

MQZ_INQUIRE

Queries the supported functionality of the authorization service.

MQZ_REFRESH_CACHE

Refresh all authorizations.

MQZ_SET_AUTHORITY

Sets the authority that an entity has to a specified object.

MQZ_TERM_AUTHORITY

Terminates authorization service component.

These entry points support the use of the Windows Security Identifier (NT SID).

These names are defined as **typedef** s, in the header file `cmqzc.h`, which can be used to prototype the component functions.

The initialization function (**MQZ_INIT_AUTHORITY**) must be the main entry point for the component. The other functions are invoked through the entry point address that the initialization function has added into the component entry point vector.

See “Creating your own service component on IBM i” on page 329 for more information.

Writing and compiling API exits

API exits let you write code that changes the behavior of IBM MQ API calls, such as MQPUT and MQGET, and then insert that code immediately before or immediately after those calls.

Note: Not supported on IBM MQ for z/OS.

Why use API exits?

Each of your applications has a specific job to do, and its code should do that task as efficiently as possible. At a higher level, you might want to apply standards or business processes to a particular queue manager for **all** the applications that use that queue manager. It is more efficient to do this above the level of individual applications, and thus without having to change the code of each application affected.

Here are a few suggestions of areas in which API exits might be useful:

- For *security*, you can provide authentication, checking that applications are authorized to access a queue or queue manager. You can also police applications' use of the API, authenticating the individual API calls, or even the parameters they use.
- For *flexibility*, you can respond to rapid changes in your business environment without changing the applications that rely on the data in that environment. You could, for example, have API exits that respond to changes in interest rates, currency exchange rates, or the price of components in a manufacturing environment.
- For *monitoring* use of a queue or queue manager, you can trace the flow of applications and messages, log errors in the API calls, set up audit trails for accounting purposes, or collect usage statistics for planning purposes.

What happens when an API exit runs?

Once you have written an exit program and identified it to IBM MQ, the queue manager automatically invokes your exit code at the registered points.

The API exit routines to run are identified in stanzas on IBM i , Windows, UNIX and Linux systems. This topic covers the stanzas in the configuration files `mq5.ini` and `qm.ini`.

The definition of the routines can occur in three places:

1. `ApiExitCommon`, in the `mqs.ini` file, identifies routines, for the whole of IBM MQ, applied when queue managers start. These can be overridden by routines defined for individual queue managers (see item 3 in this list).
2. `ApiExitTemplate`, in the `mqs.ini` file, identifies routines, for the whole of IBM MQ, copied to the `ApiExitLocal` set (see item 3 in this list) when a new queue manager is created.
3. `ApiExitLocal`, in the `qm.ini` file, identifies routines that apply to a particular queue manager.

When a new queue manager is created, the `ApiExitTemplate` definitions in `mqs.ini` are copied to the `ApiExitLocal` definitions in `qm.ini` for the new queue manager. When a queue manager is started, both the `ApiExitCommon` and `ApiExitLocal` definitions are used. The `ApiExitLocal` definitions replace the `ApiExitCommon` definitions if both identify a routine of the same name. The `Sequence` attribute, described in “Configuring API exits” on page 339 determines the order in which the routines defined in the stanzas run.

Using API exits across multiple installations of IBM MQ

Ensure that the API exits written for the earlier version of IBM MQ are used to work with all versions because the changes made to exits in Version 7.1 might not work with an earlier version. For more information about the changes made to exits, see “Writing exits and installable services on UNIX, Linux and Windows” on page 311.

The samples provided for API exits `amqsaem` and `amqsaxe` reflect the changes required while writing exits. The client application must ensure that the correct IBM MQ libraries that correspond to the installation of the queue manager with which the application is associated are linked to it prior to the launch of the application.

Writing API exits:

You can write exits for every API call using the C programming language.

Exits are available for every API call, as follows:

- `MQCB`, to reregister a callback for the specified object handle and control activation and changes to the callback
- `MQCTL`, to perform controlling actions on the object handles opened for a connection
- `MQCONN/MQCONNX`, to provide a queue manager connection handle for use on subsequent API calls
- `MQDISC`, to disconnect from a queue manager
- `MQBEGIN`, to begin a global unit of work (UOW)
- `MQBACK`, to back out a UOW
- `MQCMIT`, to commit a UOW
- `MQOPEN`, to open an IBM MQ resource for subsequent access
- `MQCLOSE`, to close an IBM MQ resource that had previously been opened for access
- `MQGET`, to retrieve a message from a queue that has previously been opened for access
- `MQPUT1`, to place a message on to a queue
- `MQPUT`, to place a message on to a queue that has previously been opened for access
- `MQINQ`, to inquire on the attributes of an IBM MQ resource that has previously been opened for access
- `MQSET`, to set the attributes of a queue that has previously been opened for access
- `MQSTAT`, to retrieve status information
- `MQSUB`, to register the applications subscription to a particular topic
- `MQSUBRQ`, to make a request for a subscription

Within API exits, the calls take the general form:

`MQ_call_EXIT (parameters, context, ApiCallParameters)`

where *call* is the MQI call name without the MQ prefix; for example, PUT, GET. The *parameters* control the function of the exit, primarily providing communication between the exit and the external control blocks MQAXP (the API exit parameter structure) and MQAXC (the API exit context structure). *context* describes the context in which the API exit was called, and *ApiCallParameters* represent the parameters to the MQI call.

To help you write your API exit, a sample exit, `amqsaxe0.c`, is provided; this exit generates trace entries to a file that you specify. You can use this sample as your starting point when writing exits.

For more information about using the sample exit, see “The API exit sample program” on page 475. For more information about the API exit calls, external control blocks, and associated topics, see API exit reference.

For general information on how to write, compile and configure an exit, see “Writing exits and installable services on UNIX, Linux and Windows” on page 311.

Using message handles in API exits

You can control which message properties an API exit has access to. Properties are associated with an `ExitMsgHandle`. Properties set in a put exit are set on the message being put, but properties retrieved in a get exit are not returned to the application.

When you register an `MQ_INIT_EXIT` exit function using the `MQXEP` MQI call with **Function** set to `MQXF_INIT` and **ExitReason** set to `MQXR_CONNECTION`, you pass in an `MQXEPO` structure as the **ExitOpts** parameter. The `MQXEPO` structure contains the `ExitProperties` field, which specifies the set of properties to be made available to the exit. It is specified as a character string representing the prefix of the properties, which corresponds to an `MQRFH2` folder name.

Each API exit receives an `MQAXP` structure, containing an `ExitMsgHandle` field. This field is set to a value generated by IBM MQ and is specific to a connection. The handle is therefore unchanged between API exits of the same or different types on the same connection.

In an `MQ_PUT_EXIT` or `MQ_PUT1_EXIT` with an **ExitReason** of `MQXR_BEFORE`, that is, an API exit performed before putting a message, any properties (other than message descriptor properties) associated with the `ExitMsgHandle` when the exit completes are set on the message being put. To prevent this happening, set `ExitMsgHandle` to `MQHM_NONE`. You can also supply a different message handle.

In an `MQ_GET_EXIT`, the `ExitMsgHandle` is cleared of properties and populated with the properties specified in the `ExitProperties` field when the `MQ_INIT_EXIT` was registered, other than message descriptor properties. These properties are not made available to the getting application. If the getting application specified a message handle in the `MQGMO` (Get message options) field, then any properties associated with that handle, including message descriptor properties, are available to the API exit. To prevent the `ExitMsgHandle` being populated with properties, set it to `MQHM_NONE`.

A sample program, `amqsaem0.c`, is provided to illustrate the use of message handles in API exits.

Compiling API exits:

After you have written an exit, you compile and link it as follows.


The following examples show the commands used for the sample program described in “The API exit sample program” on page 475. For platforms other than Windows systems, you can find the sample API exit code in `MQ_INSTALLATION_PATH/samp` and the compiled and linked shared library in `MQ_INSTALLATION_PATH/samp/bin`. For Windows systems, you can find the sample API exit code in `MQ_INSTALLATION_PATH \Tools\c\Samples`. `MQ_INSTALLATION_PATH` represents the directory in which IBM MQ was installed.

Note to users:

1. Guidance on programming 64 bit applications is listed in Coding standards on 64-bit platforms

With the introduction of Multicast clients, API exits and data-conversion exits need to be able to run on the client-side because some messages might not go through the queue manager. The following libraries are now part of the client packages as well as the server packages:

Table 42. Libraries that are now in the client and server packages

Operating system	Libraries
Windows	32 bit & 64 bit: mqm.dll & mqm.pdb
Linux & HP-UX	32 bit & 64 bit: libmqm.so & libmqm_r.so
AIX	32 bit & 64 bit: libmqm.a & libmqm_r.a
Solaris	32 bit & 64 bit: libmqm.so
 IBM i	LIBMQM & LIBMQM_R

Compiling API exits on Unix and Linux systems:

Examples of how to Compile API exits on UNIX and Linux systems.

On all platforms, the entry point to the module is MQStart.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

On AIX

Compile the API exit source code by issuing one of the following commands:

32 bit applications

Non-threaded

```
cc -e MQStart -bE:amqsaxe.exp -bM:SRE -o /var/mqm/exits/amqsaxe \
amqsaxe0.c -I MQ_INSTALLATION_PATH/inc
```

Threaded

```
xlc_r -e MQStart -bE:amqsaxe.exp -bM:SRE -o /var/mqm/exits/amqsaxe_r \
amqsaxe0.c -I MQ_INSTALLATION_PATH/inc
```

64 bit applications

Non-threaded

```
cc -q64 -e MQStart -bE:amqsaxe.exp -bM:SRE -o /var/mqm/exits64/amqsaxe \
amqsaxe0.c -I MQ_INSTALLATION_PATH/inc
```

Threaded

```
xlc_r -q64 -e MQStart -bE:amqsaxe.exp -bM:SRE -o /var/mqm/exits64/amqsaxe_r \
amqsaxe0.c -I MQ_INSTALLATION_PATH/inc
```

On HP-UX Itanium platform

32 bit applications

Non-threaded

Compile the API Exit source code:

```
c89 +e +z -c -D_HPUX_SOURCE -o amqsaxe.o amqsaxe0.c -I MQ_INSTALLATION_PATH/inc
```

Link the API Exit source code

```
ld +b: -b amqsaxe.o +ee MQStart -o /var/mqm/exits/amqsaxe
rm amqsaxe.o
```

Threaded

Compile the API Exit source code:

```
c89 +e +z -c -D_HPUX_SOURCE -o amqsaxe.o amqsaxe0.c -I MQ_INSTALLATION_PATH/inc
```

Link the API Exit source code

```
ld +b: -b amqsaxe.o +ee MQStart -o /var/mqm/exits/amqsaxe_r
rm amqsaxe.o
```

64 bit applications

Non-threaded

Compile the API Exit source code:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o amqsaxe.o amqsaxe0.c -I MQ_INSTALLATION_PATH/inc
```

Link the API Exit source code

```
ld -b amqsaxe.o +ee MQStart -o /var/mqm/exits64/amqsaxe
rm amqsaxe.o
```

Threaded

Compile the API Exit source code:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o amqsaxe.o amqsaxe0.c -I MQ_INSTALLATION_PATH/inc
```

Link the API Exit source code

```
ld -b amqsaxe.o +ee MQStart -o /var/mqm/exits64/amqsaxe_r
rm amqsaxe.o
```

On Linux

Compile the API exit source code by issuing one of the following commands:

31 bit applications

Non-threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/amqsaxe amqsaxe0.c \
-I MQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/amqsaxe_r amqsaxe0.c \
-I MQ_INSTALLATION_PATH/inc
```

32 bit applications

Non-threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/amqsaxe amqsaxe0.c \
-I MQ_INSTALLATION_PATH/inc
```

Threaded


```
gcc -m32 -shared -fPIC -o /var/mqm/exits/amqsaxe_r amqsaxe0.c \  
-I MQ_INSTALLATION_PATH/inc
```

64 bit applications

Non-threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/amqsaxe amqsaxe0.c \  
-I MQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/amqsaxe_r amqsaxe0.c \  
-I MQ_INSTALLATION_PATH/inc
```

On Solaris

Compile the API exit source code by issuing one of the following commands:

32 bit applications

SPARC platform

```
cc -xarch=v8plus -KPIC -mt -G -o /var/mqm/exits/amqsaxe \  
amqsaxe0.c -I MQ_INSTALLATION_PATH/inc \  
-R/usr/lib/32 -lsocket -lnsl -ldl
```

x86-64 platform

```
cc -xarch=386 -KPIC -mt -G -o /var/mqm/exits/amqsaxe \  
amqsaxe0.c -I MQ_INSTALLATION_PATH/inc \  
-R/usr/lib/32 -lsocket -lnsl -ldl
```

64 bit applications

SPARC platform

```
cc -xarch=v9 -KPIC -mt -G -o /var/mqm/exits64/amqsaxe \  
amqsaxe0.c -I MQ_INSTALLATION_PATH/inc \  
-R/usr/lib/64 -lsocket -lnsl -ldl
```

x86-64 platform

```
cc -xarch=amd64 -KPIC -mt -G -o /var/mqm/exits64/amqsaxe \  
amqsaxe0.c -I MQ_INSTALLATION_PATH/inc \  
-R/usr/lib/64 -lsocket -lnsl -ldl
```

On Windows systems:

Compile and link the sample API exit program, `amqsaxe0.c`, on Windows

A manifest file is an optional XML document containing the version, or any other, information that can be embedded in a compiled application or DLL.

If you have no such document, omit the `-manifest manifest.file` parameter in the `mt` command.

Adapt the commands in the examples in Figure 56 on page 338 or Figure 57 on page 338 to compile and link `amqsaxe0.c` on Windows. The commands work with Microsoft Visual Studio 2008, 2010, or 2012. The examples assume that the `C:\Program Files\IBM\MQ\tools\c\samples` directory is the current directory.

32 bit

```
cl /c /nologo /MD /Foamsaxe0.obj amqsaxe0.c

link /nologo /dll /def:amsaxe.def

amsaxe0.obj \
  /manifest /out:amsaxe.dll

mt -nologo -manifest amsaxe.dll.manifest \
  -outputresource:amsaxe.dll;2
```

Figure 56. Compile and link amqsaxe0.c on 32 bit Windows

64 bit

```
cl /c /nologo /MD /Foamsaxe0.obj amqsaxe0.c

link /nologo /dll /def:amsaxe.def \
  /libpath:..\..\lib64 \

amsaxe0.obj /manifest /out:amsaxe.dll

mt -nologo -manifest amsaxe.dll.manifest \
  -outputresource:amsaxe.dll;2
```

Figure 57. Compile and link amqsaxe0.c on 64 bit Windows

Related concepts:

“The API exit sample program” on page 475

The sample API exit generates an MQI trace to a user-specified file with a prefix defined in the MQAPI_TRACE_LOGFILE environment variable.

On IBM i:

Compiling API exits on IBM i.


An exit is created as follows (for a C language example):


1. Create a module using CRTCMOD. Compile it to use teraspace by including the parameter TERASPACE(*YES *TSIFC).
2. Create a service program from the module using CRTSRVPGM. You must bind it to the service program QMQM/LIBMQMZF_R for multithreaded API exits.

Configuring API exits:

You configure IBM MQ to enable API exits by changing the configuration information.

To change the configuration information, you must change the stanzas that define the exit routines and the sequence in which they run. This information can be changed in the following ways:

- Using the IBM MQ Explorer (On Windows and Linux (x86 and x86-64 platforms))
- Using the **amqmdain** command (On Windows)
- Using the mqs.ini and qm.ini files directly (On Windows,  IBM i, UNIX and Linux systems).

The mqs.ini file contains information relevant to all the queue managers on a particular node. You can find it in the /var/mqm directory on UNIX and Linux , in the /QIBM/UserData/mqm directory on IBM i and in the WorkPath specified in the HKLM\SOFTWARE\IBM\WebSphere MQ key on Windows systems.

The qm.ini file contains information relevant to a specific queue manager. There is one queue manager configuration file for each queue manager, 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:

On UNIX and Linux systems:

```
/var/mqm/qmgrs/QMNAME/qm.ini
```

 On IBM i systems:

```
/QIBM/UserData/mqm/qmgrs/QMNAME/qm.ini
```

On Windows systems:

```
C:\ProgramData\IBM\MQ\qmgrs\QMNAME\qm.ini
```

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

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.)

Stanzas to configure

The stanzas that must be changed are the following:

ApiExitCommon

Defined in mqs.ini and in the IBM MQ Explorer on the IBM MQ properties page, under Exits. When any queue manager starts, the attributes in this stanza are read, and then overridden by the API exits defined in qm.ini.

ApiExitTemplate

Defined in mqs.ini and in the IBM MQ Explorer on the IBM MQ properties page, under Exits. When any queue manager is created, the attributes in this stanza are copied into the newly created qm.ini file under the ApiExitLocal stanza.

ApiExitLocal

Defined in qm.ini and in the IBM MQ Explorer on the queue manager properties page, under Exits. When the queue manager starts, API exits defined here override the defaults defined in mqs.ini.

Attributes for the stanzas

- Name the API exit using the following attribute:

Name=ApiExit_name

The descriptive name of the API exit passed to it in the ExitInfoName field of the MQAXP structure.

This name must be unique, no longer than 48 characters, and contain only valid characters for the names of IBM MQ objects (for example, queue names).

- Identify the module and entry point of the API exit code to run using the following attributes:

Function=function_name

The name of the function entry point into the module containing the API exit code. This entry point is the MQ_INIT_EXIT function.

The length of this field is limited to MQ_EXIT_NAME_LENGTH.

Module=module_name

The module containing the API exit code.

If this field contains the full path name of the module it is used as is.

If this field contains just the module name, the module is located using the ExitsDefaultPath attribute in the ExitPath in qm.ini.

On platforms that support separate threaded libraries, you must provide both a non-threaded and a threaded version of the API exit module. The threaded version must have an `_r` suffix. The threaded version of the IBM MQ application stub implicitly appends `_r` to the given module name before it is loaded.

The length of this field is limited to the maximum path length the platform supports.

- Optionally pass data with the exit using the following attribute:

Data=data_name

Data to be passed to the API exit in the ExitData field of the MQAXP structure.

If you include this attribute, leading and trailing blanks are removed, the remaining string is truncated to 32 characters, and the result is passed to the exit. If you omit this attribute, the default value of 32 blanks is passed to the exit.

The maximum length of this field is 32 characters.

- Identify the sequence of this exit in relation to other exits using the following attribute:

Sequence=sequence_number

The sequence in which this API exit is called relative to other API 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 contiguous. A sequence of 1, 2, 3 has the same result as a sequence of 7, 42, 1096. If two exits have the same sequence number, the queue manager decides which one to call first. You can tell which was called after the event by putting the time or a marker in ExitChainArea indicated by the ExitChainAreaPtr in MQAXP or by writing your own log file.

This attribute is an unsigned numeric value.

Sample stanzas

The sample mqs.ini file contains the following stanzas:

ApiExitTemplate

This stanza defines an exit with the descriptive name `OurPayrollQueueAuditor`, module name `auditor`, and sequence number 2. A data value of 123 is passed to the exit.

ApiExitCommon

This stanza defines an exit with the descriptive name MQPoliceman, module name tmqp, and sequence number 1. The data passed is an instruction (CheckEverything).

mqs.ini

```
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
```

The following sample qm.ini file contains an ApiExitLocal definition of an exit with the descriptive name ClientApplicationAPIchecker, module name ClientAppChecker, and sequence number 3.

qm.ini

```
ApiExitLocal:
  Name=ClientApplicationAPIchecker
  Sequence=3
  Function=EntryPoint
  Module=/usr/Dev/ClientAppChecker
  Data=9.20.176.20
```

Channel-exit programs for messaging channels

This collection of topics contains information about IBM MQ channel-exit programs for messaging channels.

Message channel agents (MCAs) can also call data-conversion exits. For more information about writing data-conversion exits, see “Writing data-conversion exits” on page 364.

Some of this information also applies to exits on MQI channels, which connect IBM MQ MQI clients to queue managers. For more information, see Channel-exit programs for MQI channels.

Channel-exit programs are called at defined places in the processing carried out by MCA programs.

Some of these user-exit programs work in complementary pairs. For example, if a user-exit program is called by the sending MCA to encrypt the messages for transmission, the complementary process must be functioning at the receiving end to reverse the process.

Table 43 shows the types of channel exit that are available for each channel type.

Table 43. Channel exits available for each channel type

Channel Type	Message exit	Message- retry exit	Receive exit	Security exit	Send exit	Auto-definition exit
Sender channel	Yes		Yes	Yes	Yes	
Server channel	Yes		Yes	Yes	Yes	
Cluster- sender channel	Yes		Yes	Yes	Yes	Yes
Receiver channel	Yes	Yes	Yes	Yes	Yes	Yes

Table 43. Channel exits available for each channel type (continued)

Channel Type	Message exit	Message- retry exit	Receive exit	Security exit	Send exit	Auto-definition exit
Requester channel	Yes	Yes	Yes	Yes	Yes	
Cluster-receiver channel	Yes	Yes	Yes	Yes	Yes	Yes
Client-connection channel			Yes	Yes	Yes	
Server-connection channel			Yes	Yes	Yes	Yes

Notes: z/OS

1. On z/OS, the auto-definition exit applies to cluster-sender and cluster-receiver channels only.

If you are going to run channel exits on a client, you cannot use the MQSERVER environment variable. Instead, create and reference a client channel definition table (CCDT) as described in Client channel definition table.

Processing overview:

An overview of how MCAs use channel-exit programs.

On startup, the MCAs exchange a startup dialog to synchronize processing. Then they switch to a data exchange that includes the security exits. These exits must end successfully for the startup phase to complete and to allow messages to be transferred.

The security check phase is a loop, as shown in Figure 58.

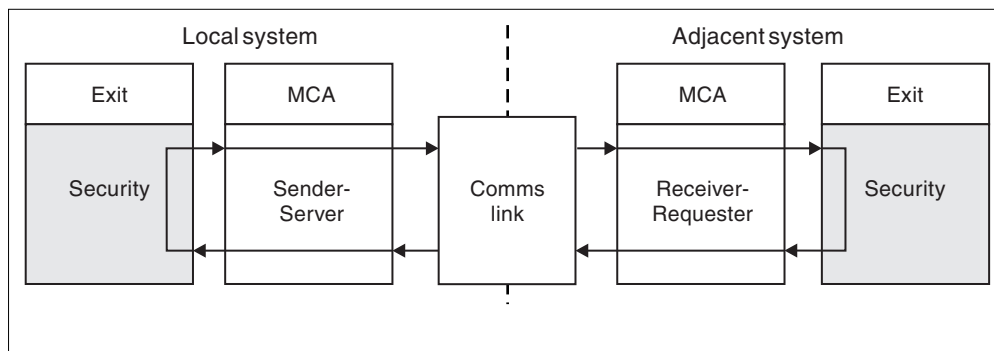


Figure 58. Security exit loop

During the message transfer phase, the sending MCA gets messages from a transmission queue, calls the message exit, calls the send exit, and then sends the message to the receiving MCA, as shown in Figure 59 on page 343.

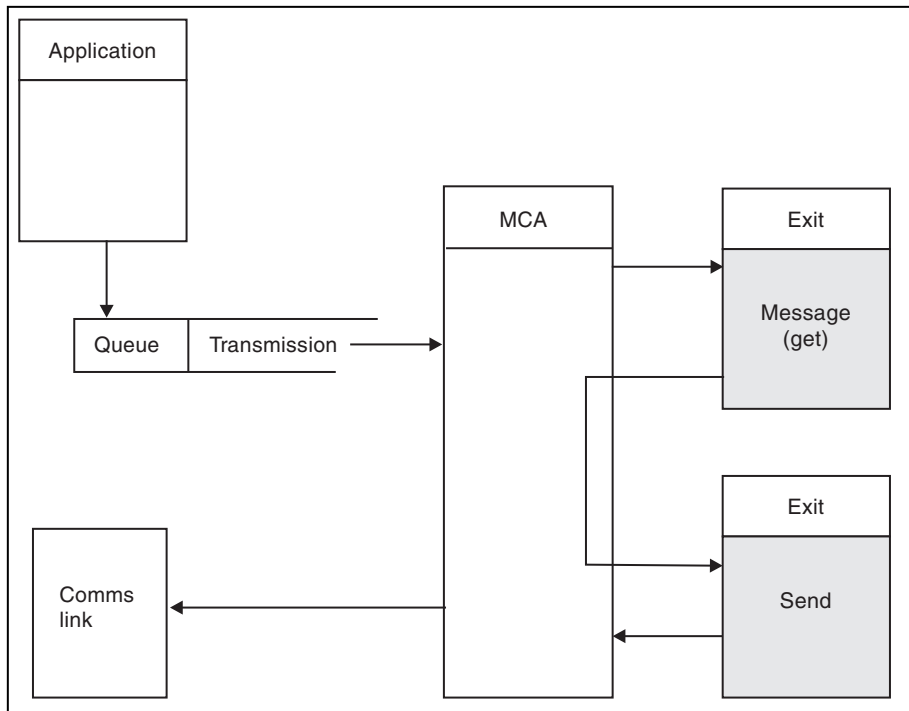


Figure 59. Example of a send exit at the sender end of message channel

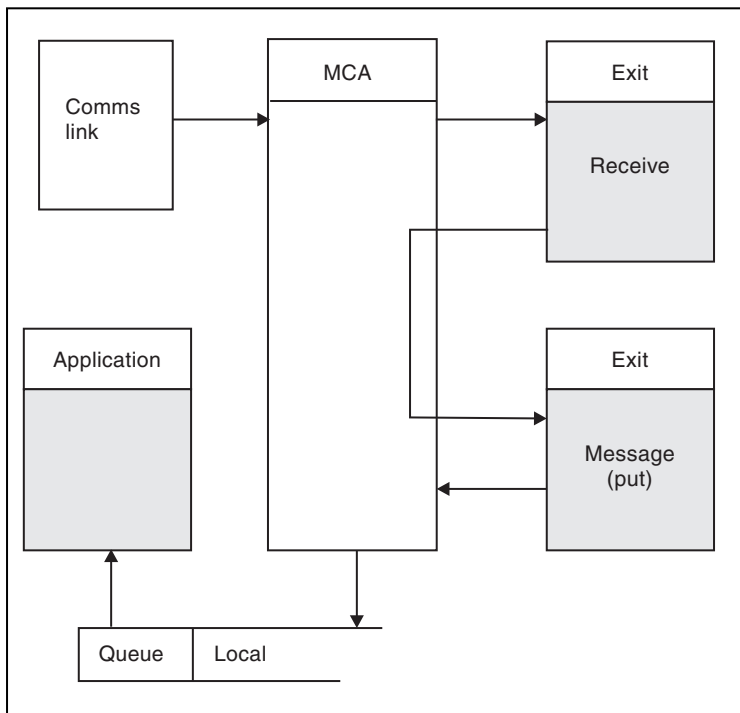


Figure 60. Example of a receive exit at the receiver end of message channel

The receiving MCA receives a message from the communications link, calls the receive exit, calls the message exit, and then puts the message on the local queue, as shown in Figure 60. (The receive exit can be called more than once before the message exit is called.)

Writing channel-exit programs:

You can use the following information to help you write channel-exit programs.

User exits and channel-exit programs can use all MQI calls, except as noted in the sections that follow. For MQ V7 and later, the MQCXP structure version 7 and higher contains the connection handle hConn, which can be used instead of issuing MQCONN. For earlier versions, to obtain the connection handle, an MQCONN must be issued, even though an MQRC_ALREADY_CONNECTED warning is returned because the channel itself is connected to the queue manager.

Note that the channel exit must be threadsafe.

For exits on client-connection channels, the queue manager to which the exit tries to connect depends on how the exit was linked. If the exit was linked with MQM.LIB (or QMQM/LIBMQM on IBM i) and you do not specify a queue manager name on the MQCONN call, the exit tries to connect to the default queue manager on your system. If the exit was linked with MQM.LIB (or QMQM/LIBMQM on IBM i) and you specify the name of the queue manager that was passed to the exit through the QMgrName field of MQCD, the exit tries to connect to that queue manager. If the exit was linked with MQIC.LIB or any other library, the MQCONN call fails whether you specify a queue manager name or not.

You should avoid altering the state of the transaction associated with the passed hConn in a channel exit; you must not use the MQCMIT, MQBACK or MQDISC verbs with the channel hConn, and you cannot use the MQBEGIN verb specifying the channel hConn.



If MQCONNX is used specifying MQCNO_HANDLE_SHARE_BLOCK or MQCNO_HANDLE_SHARE_NO_BLOCK to create a new IBM MQ connection, then it is your responsibility to ensure that the connection is correctly managed and disconnects from the queue manager correctly. For example, a channel exit that creates a new connection to the queue manager on every invocation without disconnecting, results in connection handles building up and an increase in the number of agent threads.

An exit runs in the same thread as the MCA itself and uses the same connection handle. So, it runs inside the same UOW as the MCA and any calls made under sync point are committed or backed out by the channel at the end of the batch.

Therefore, a channel message exit could send notification messages that are only committed to that queue when the batch containing the original message is committed. So, it is possible to issue sync point MQI calls from a channel message exit.

A channel exit can change fields in the MQCD. However, these changes are not acted on, except in the circumstances listed. If a channel exit program changes a field in the MQCD data structure, the new value is ignored by the IBM MQ channel process. However, the new value remains in the MQCD and is passed to any remaining exits in an exit chain and to any conversation sharing the channel instance. For more information, see [Changing MQCD fields in a channel exit](#)

Also, for programs written in C, non-reentrant C library function must not be used in a channel-exit program.

  If you use multiple channel exit libraries simultaneously, problems can arise on some UNIX and Linux platforms if the code for two different exits contains identically named functions. When a channel exit is loaded, the dynamic loader resolves function names in the exit library to the addresses where the library is loaded. If two exit libraries define separate functions which happen to have identical names, this resolution process might incorrectly resolve the function names of one library to use the functions of another. If this problem occurs, specify to the linker that it must only

export the required exit and MQStart functions, as these functions are unaffected. Other functions must be given local visibility so that they are not used by functions outside their own exit library. Consult the documentation for the linker for more information.

All exits are called with a channel exit parameter structure (MQCXP), a channel definition structure (MQCD), a prepared data buffer, data length parameter, and buffer length parameter. The buffer length must not be exceeded:

- For message exits, you must allow for the largest message required to be sent across the channel, plus the length of the MQXQH structure.
- For send and receive exits, the largest buffer you must allow for is as follows:

LU 6.2 32 KB

TCP:  IBM i 16 KB

 Others 32 KB

Note: The maximum usable length might be 2 bytes less than this length. Check the value returned in MaxSegmentLength for details. For more information about MaxSegmentLength, see MaxSegmentLength.

NetBIOS:


64 KB

SPX: 64 KB

Note: Receive exits on sender channels and sender exits on receiver channels use 2 KB buffers for TCP.

- For security exits, the distributed queuing facility allocates a buffer of 4000 bytes.

It is permissible for the exit to return an alternative buffer, together with the relevant parameters. See “Channel-exit programs for messaging channels” on page 341 for call details.

Writing channel exit programs on z/OS: 

You can use the following information to help you write and compile channel-exit programs for z/OS.

The exits are started as if by a z/OS LINK, in:

- Non-authorized problem program state
- Primary address space control mode
- Non-cross-memory mode
- Non-access register mode
- 31 bit addressing mode

The link-edited modules must be placed in the data set specified by the CSQXLIB DD statement of the channel initiator address space procedure; the names of the load modules are specified as the exit names in the channel definition.

When writing channel exits for z/OS, the following rules apply:

- Exits must be written in assembler or C; if C is used, it must conform to the C systems programming environment for system exits, described in the z/OS C/C++ Programming Guide.
- Exits are loaded from the non-authorized libraries defined by a CSQXLIB DD statement. Providing CSQXLIB has DISP=SHR, exits can be updated while the channel initiator is running. The new version is used when the channel is restarted.
- Exits must be reentrant, and capable of running anywhere in virtual storage.

- Exits must reset the environment, on return, to that at entry.
- Exits must free any storage obtained, or ensure that it is freed by a subsequent exit invocation.
For storage that is to persist between invocations, use the z/OS STORAGE service, or the 4kmalc library function for System Programming C.
For more information about this function, see 4kmalc() -- Allocate Page-Aligned Storage.
- All IBM MQ MQI calls except MQCMIT or CSQBCMT and MQBACK or CSQBBAK can be used. They must be contained after MQCONN (with a blank queue manager name). If these calls are used, the exit must be link-edited with the stub CSQXSTUB.

The exception to this rule is that security channel exits can issue commit and backout MQI calls. To issue such calls, code the verbs CSQXCMT and CSQXBAK in place of MQCMIT or CSQBCMT and MQBACK or CSQBBAK.

- All exits that use stub CSQXSTUB from IBM WebSphere MQ Version 7.0 or later must be link-edited in a CSQXLIB load library with format PDS-E.
- Exits must not use any system services that cause a wait, because using system services would severely affect the handling of some or all the other channels. Many channels are run under a single TCB typically. If you do something in an exit that causes a wait and you do not use MQXWAIT, it causes all these channels to wait. Causing channels to wait does not give any functional problems, but might have an adverse effect on performance. Most SVCs involve waits, so you must avoid them, except for the following SVCs:
 - GETMAIN/FREEMAIN/STORAGE
 - LOAD/DELETE

In general, therefore, avoid SVCs, PCs, and I/O. Instead, use the MQXWAIT call.

- Exits do not issue ESTAEs or SPIEs, apart from in any subtasks they attach, because their error handling might interfere with the error handling performed by IBM MQ. This means that IBM MQ might not be able to recover from an error, or that your exit program might not receive all the error information.
- The MQXWAIT call (see MQXWAIT) provides a wait service that waits for I/O and other events; if this service is used, exits must not use the linkage stack.
For I/O and other facilities that do not provide non-blocking facilities or an ECB to wait on, a separate subtask must be ATTACHED, and its completion waited for by MQXWAIT; because of the processing that this technique incurs, this facility must be used only by the security exit.
- The MQDISC MQI call does not cause an implicit commit to occur within the exit program. A commit of the channel process is performed only when the channel protocol dictates.

The following exit samples are provided with IBM MQ for z/OS:

CSQ4BAX0

This sample is written in assembler, and illustrates the use of MQXWAIT.

CSQ4BCX1 and CSQ4BCX2

These samples are written in C and illustrate how to access the parameters.

CSQ4BCX3 and CSQ4BAX3

These samples are written in C and assembler respectively.

The CSQ4BCX3 sample (which is pre-compiled into the SCSQAUTH LOADLIB, should function with no changes necessary on the exit itself. You can create a LOADLIB (for example, called MY.TEST.LOADLIB) and copy the SCSQAUTH(CSQ4BCX3) member to it.

To set up a security exit on a client connection, carry out the following procedure:

1. Establish a valid OMVS segment for the user ID that the channel initiator uses.
This allows the IBM MQ for z/OS channel initiator to use TCP/IP with the UNIX System Services (USS) socket interface, in order to facilitate exit processing. Note that it is unnecessary to define an OMVS segment for the user ID of any connecting client.

2. Ensure that the exit code itself runs only in a program controlled environment.
This means everything loaded into the CHINIT address space must be loaded from a program controlled library (meaning all libraries in the STEPLIB), and any libraries named on CSQXLIB and

```
++h1q++.SCSQANLx
++h1q++.SCSQMVR1
++h1q++.SCSQAUTH
```

To set a load library as program controlled, use a command similar to this example:

```
RALTER PROGRAM * ADDMEM('MY.TEST.LOADLIB' //NOPADCHK)
```

Then you can activate or refresh the program controlled environment by issuing the command:

```
SETROPTS WHEN(PROGRAM) REFRESH
```

3. Add the exit LOADLIB to the CSQXLIB DD (in the CHINIT started procedure), by issuing the following command:

```
ALTER CHANNEL(xxxx) CHLTYPE(SVRCONN)SCYEXIT(CSQ4BCX3)
```

This activates the exit for the named channel.


4. Your external security manager (ESM) lists any other libraries to be program controlled, but note that none of the ESM or C libraries needs to be under program control.

See IBM MQ for z/OS server connection channel for further information on setting up a security exit using the sample CSQ4BCX3.


CSQ4BCX4

This sample is written in C and demonstrates using the **RemoteProduct** and **RemoteVersion** fields in MQCXP.

Related concepts:

 "Writing channel exit programs on IBM i"


You can use the following information to help you write and compile channel-exit programs for IBM i.

 "Writing channel-exit programs on UNIX, Linux, and Windows" on page 348

You can use the following information to help you write channel-exit programs for UNIX, Linux, and Windows systems.

Related information:

IBM MQ for z/OS server connection channel

Writing channel exit programs on IBM i: 

You can use the following information to help you write and compile channel-exit programs for IBM i.

The exit is a program object written in the ILE C, ILE RPG, or ILE COBOL language. The exit program names and their libraries are named in the channel definition.

Observe the following conditions when creating and compiling an exit program:

- The program must be made thread safe and created with the ILE C, ILE RPG, or ILE COBOL compiler. For ILE RPG you must specify the THREAD(*SERIALIZE) control specification, and for ILE COBOL you must specify SERIALIZE for the THREAD option of the PROCESS statement. The programs must also be bound to the threaded IBM MQ libraries: QMQM/LIBMQM_R in the case of ILE C and ILE RPG, and AMQOSTUB_R in the case of ILE COBOL. For additional information about making RPG or COBOL applications thread safe, refer to the appropriate Programmer's Guide for the language.
- IBM MQ for IBM i requires that the exit programs are enabled for teraspace support. (Teraspace is a form of shared memory introduced in OS/400 V4R4.) For the ILE RPG and COBOL compilers, any


programs compiled on OS/400 V4R4 or later are so enabled. For C, the programs must be compiled with the TERASPACE(*YES *TSIFC) options specified on CRTCMOD or CRTBNDC commands.

- An exit returning a pointer to its own buffer space must ensure that the object pointed to exists beyond the time span of the channel-exit program. The pointer cannot be the address of a variable on the program stack, nor of a variable in the program heap. Instead, the pointer must be obtained from the system. An example is a user space created in the user exit. To ensure that any data area allocated by the channel-exit program is still available for the MCA when the program ends, the channel exit must run in the activation group of the caller or a named activation group. Do this by setting the ACTGRP parameter on CRTPGM to a user-defined value or *CALLER. If the program is created in this way, the channel-exit program can allocate dynamic memory and pass a pointer to this memory back to the MCA.

Related concepts:

 [“Writing channel-exit programs on UNIX, Linux, and Windows”](#)

You can use the following information to help you write channel-exit programs for UNIX, Linux, and Windows systems.

 [“Writing channel exit programs on z/OS”](#) on page 345

You can use the following information to help you write and compile channel-exit programs for z/OS.

Writing channel-exit programs on UNIX, Linux, and Windows: 

You can use the following information to help you write channel-exit programs for UNIX, Linux, and Windows systems.

Follow the instructions outlined in “Writing exits and installable services on UNIX, Linux and Windows” on page 311. Use the following channel exit specific information, where appropriate:

The exit must be written in C, and is a DLL on Windows.

Define a dummy MQStart() routine in the exit and specify MQStart as the entry point in the library. Figure 61 shows how to set up an entry to your program:

```
#include <cmqec.h>


void MQStart() {;} /* dummy entry point - for consistency only */
void MQENTRY ChannelExit ( PMQEXP pChannelExitParms,
                           PMQCD  pChannelDefinition,
                           PMQLONG pDataLength,
                           PMQLONG pAgentBufferLength,
                           PMQVOID pAgentBuffer,
                           PMQLONG pExitBufferLength,
                           PMQPTR  pExitBufferAddr)
{
  ... Insert code here
}
```

Figure 61. Sample source code for a channel exit


When writing channel exits for Windows using Visual C++, you must write your own DEF file. An example of how is shown in Figure 62 on page 349. For further information on writing channel exit programs, see “Writing channel-exit programs” on page 344.

Figure 62. Sample DEF file for Windows

Related concepts:

 “Writing channel exit programs on IBM i” on page 347

You can use the following information to help you write and compile channel-exit programs for IBM i.

 “Writing channel exit programs on z/OS” on page 345

You can use the following information to help you write and compile channel-exit programs for z/OS.

Channel security exit programs:

You can use security exit programs to verify that the partner at the other end of a channel is genuine. This is known as authentication. To specify that a channel must use a security exit, specify the exit name in the SCYEXIT field of the channel definition.

Note: Authentication can also be achieved with channel authentication records. Channel authentication records provide great flexibility in preventing access to queue managers from certain users and channels, and in mapping remote users to IBM MQ user identifiers. TLS support is also provided by IBM MQ to authenticate your users and to provide encryption and data integrity checks for your data. For more information about TLS, see TLS security protocols in IBM MQ. However, if you still require more sophisticated (or different) forms of security processing, and other types of checks and security context establishment, consider writing security exits.

For security exits written prior to IBM WebSphere MQ Version 7.1 it is worth noting that earlier versions of IBM MQ queried the underlying secure sockets provider (e.g. GSKit) to determine the remote partner's certificate Subject Distinguished Name (SSLPEER) and Issuer Distinguished Name (SSLCERTI). In IBM WebSphere MQ Version 7.1 support was added for a range of new security attributes. In order to access these attributes IBM WebSphere MQ Version 7.1 obtains the DER encoding of the certificate and uses it to determine the Subject and Issuer DN. The Subject and Issuer DN attributes appear in the following channel status attributes:

- SSLPEER (PCF selector MQCACH_SSL_SHORT_PEER_NAME)
- SSLCERTI (PCF selector MQCACH_SSL_CERT_ISSUER_NAME)

These values are returned by channel status commands as well as the data passed to channel security exits listed, as shown:

- MQCD SSLPeerNamePtr
- MQCXP SSLRemCertIssNamePtr

In IBM WebSphere MQ Version 7.1, a SERIALNUMBER attribute is also included in the Subject DN and contains the serial number for the remote partner's certificate. Also some DN attributes are returned in a different sequence from previous releases. Consequently the composition of the SSLPEER and SSLCERTI fields are altered in IBM WebSphere MQ Version 7.1 from previous releases and it is therefore recommended that any security exits or applications dependant on these fields be examined and updated.

Existing IBM MQ peer name filters specified via the SSLPEER field of a channel definition are not affected and continue to operate in the same manner as in earlier releases. This is because the IBM MQ peer name matching algorithm has been updated to process existing SSLPEER filters without any need to alter the channel definitions. This change is most likely to affect security exits and applications which depend upon the Subject DN and Issuer DN values returned by the PCF programming interface.

A security exit can be written in C or Java.

Channel security exit programs are called at the following places in the processing cycle of an MCA:

- At MCA initiation and termination.
- Immediately after the initial data negotiation is finished on channel startup. The receiver or server end of the channel can initiate a security message exchange with the remote end by providing a message to be delivered to the security exit at the remote end. It might also decline to do so. The exit program is started again to process any security message received from the remote end.
- Immediately after the initial data negotiation is finished on channel startup. The sender or requester end of the channel processes a security message received from the remote end, or initiates a security exchange when the remote end cannot. The exit program is started again to process all subsequent security messages that might be received.

A requester channel never gets called with MQXR_INIT_SEC. The channel notifies the server that it has a security exit program, and the server then has the opportunity to initiate a security exit. If it does not have one, it sends a null security flow to allow the requester to call its exit program.

Note: Avoid sending zero-length security messages.

Examples of the data exchanged by security-exit programs are illustrated in figures Figure 63 on page 351 through Figure 66 on page 354. These examples show the sequence of events that occur involving the security exit of the receiver, and the security exit of the sender. Successive rows in the figures represent the passage of time. In some cases, the events at the receiver and sender are not correlated, and therefore can occur at the same time or at different times. In other cases, an event at one exit program results in a complementary event occurring later at the other exit program. For example, in Figure 63 on page 351:

1. The receiver and sender are each invoked with MQXR_INIT, but these invocations are not correlated and can therefore occur at the same time or at different times.
2. The receiver is next invoked with MQXR_INIT_SEC, but returns MQXCC_OK which requires no complementary event at the sender exit.
3. The sender is next invoked with MQXR_INIT_SEC. This is not correlated with the invocation of the receiver with MQXR_INIT_SEC. The sender returns MQXCC_SEND_SEC_MSG, which causes a complementary event at the receiver exit.
4. The receiver is then invoked with MQXR_SEC_MSG, and returns MQXCC_SEND_SEC_MSG, which causes a complementary event at the sender exit.
5. The sender is then invoked with MQXR_SEC_MSG, and returns MQXCC_OK which requires no complementary event at the receiver exit.

Receiver exit	Sender exit
Invoked with MQXR_INIT Responds with MQXCC_OK	Invoked with MQXR_INIT Responds with MQXCC_OK
Invoked with MQXR_INIT_SEC Responds with MQXCC_OK	
	Invoked with MQXR_INIT_SEC Responds with MQXCC_SEND_SEC_MSG
Invoked with MQXR_SEC_MSG Responds with MQXCC_SEND_SEC_MSG	
<p style="text-align: center;"><i>Message transfer begins</i></p>	Invoked with MQXR_SEC_MSG Responds with MQXCC_OK

Figure 63. Sender-initiated exchange with agreement

Receiver exit	Sender exit
Invoked with MQXR_INIT Responds with MQXCC_OK	Invoked with MQXR_INIT Responds with MQXCC_OK
Invoked with MQXR_INIT_SEC Responds with MQXCC_OK	
	Invoked with MQXR_INIT_SEC Responds with MQXCC_SEND_SEC_MSG
Invoked with MQXR_SEC_MSG Responds with MQXCC_OK	
	Invoked with MQXR_SEC_MSG Responds with MQXCC_SUPPRESS_FUNCTION <i>Channel closes</i>
Invoked with MQXR_TERM Responds with MQXCC_OK	Invoked with MQXR_TERM Responds with MQXCC_OK

Figure 64. Sender-initiated exchange with no agreement

Receiver exit	Sender exit
Invoked with MQXR_INIT Responds with MQXCC_OK	Invoked with MQXR_INIT Responds with MQXCC_OK
Invoked with MQXR_INIT_SEC Responds with MQXCC_SEND_SEC_MSG	
	Invoked with MQXR_SEC_MSG Responds with MQXCC_SEND_SEC_MSG
Invoked with MQXR_SEC_MSG Responds with MQXCC_OK <i>Message transfer begins</i>	
Invoked with MQXR_TERM Responds with MQXCC_OK	Invoked with MQXR_TERM Responds with MQXCC_OK

Figure 65. Receiver-initiated exchange with agreement

Receiver exit	Sender exit
Invoked with MQXR_INIT Responds with MQXCC_OK	Invoked with MQXR_INIT Responds with MQXCC_OK
Invoked with MQXR_INIT_SEC Responds with MQXCC_SEND_SEC_MSG	
	Invoked with MQXR_SEC_MSG Responds with MQXCC_OK
Invoked with MQXR_SEC_MSG Responds with MQXCC_SUPPRESS_FUNCTION	
<i>Channel closes</i>	


Figure 66. Receiver-initiated exchange with no agreement

The channel security exit program is passed an agent buffer containing the security data, excluding any transmission headers, generated by the security exit. This data can be any suitable data so that either end of the channel is able to perform security validation.

The security exit program at both the sending and receiving end of the message channel can return either of two response codes to any call:

- Security exchange ended with no errors
- Suppress the channel and close down

Note:

1. The channel security exits typically work in pairs. When you define the appropriate channels, make sure that compatible exit programs are named for both ends of the channel.
2.  In IBM i, security exit programs that have been compiled with Use adopted authority (USEADPAUT=*YES) can adopt QMQM or QMQMADM authority. Take care that the exit does not use this feature to pose a security risk to your system.
3. On a TLS channel on which the other end of the channel provides a certificate, the security exit receives the Distinguished Name of the subject of this certificate in the MQCD field accessed by SSLPeerNamePtr and the Distinguished Name of the issuer in the MQCXP field accessed by SSLRemCertIssNamePtr. Uses to which this name can be put are:
 - To restrict access over the TLS channel.
 - To change MQCD.MCAUserIdentifier based on the name.

Related information:

Channel authentication records

Transport Layer Security (TLS) concepts

Writing a security exit:

You can write a security exit by using the security exit skeleton code.

Figure 67 illustrates how to write a security exit.

```
void MQENTRY MQStart() {;}
void MQENTRY EntryPoint (PMQVOID pChannelExitParms,
                        PMQVOID pChannelDefinition,
                        PMQLONG pDataLength,
                        PMQLONG pAgentBufferLength,
                        PMQVOID pAgentBuffer,
                        PMQLONG pExitBufferLength,
                        PMQPTR pExitBufferAddr)
{
    PMQCXP pParms = (PMQCXP)pChannelExitParms;
    PMQCD pChDef = (PMQCD)pChannelDefinition;
    /* TODO: Add Security Exit Code Here */
}
```

Figure 67. Security exit skeleton code

The standard IBM MQ Entry Point MQStart must exist, but is not required to perform any function. The name of the function (EntryPoint in this example) can be changed, but the function must be exported when the library is compiled and linked. As in the previous example, the pointers pChannelExitParms must be cast to PMQCXP and pChannelDefinition must be cast to PMQCD. For general information about calling channel exits and the use of parameters, see MQ_CHANNEL_EXIT. These parameters are used in a security exit as follows:

PMQVOID pChannelExitParms

input/output

Pointer to MQCXP structure - cast to PMQCXP to access fields. This structure is used to communicate between the Exit and MCA. The following fields in the MQCXP are of particular interest for Security Exits:

ExitReason

Tells the Security Exit the current state in the security exchange and is used when deciding what action to take.

ExitResponse

The response to the MCA which dictates the next stage in the security exchange.

ExitResponse2

Extra control flags to govern how the MCA interprets the response of the Security Exit.

ExitUserArea

16 bytes (maximum) of storage which can be used by the Security Exit to maintain state between calls.

ExitData

Contains the data specified in the SCYDATA field of the channel definition (32 bytes padded to the right with blanks).

PMQVOID pChannelDefinition

input/output

Pointer to MQCD structure - cast to PMQCD to access fields. This parameter contains the definition of the channel. The following fields in the MQCD are of particular interest for Security Exits:

ChannelName

The channel name (20 bytes padded to the right with blanks).

ChannelType

A code defining the channel type.

MCA User Identifier

This group of three fields is initialized to the value of the MCAUSER field specified in the channel definition. Any user identifier specified by the Security Exit in these fields is used for access control (not applicable to SDR, SVR, CLNTCONN, or CLUSSDR channels).

MCAUserIdentifier

First 12 bytes of identifier padded to the right with blanks.

LongMCAUserIdPtr

Pointer to a buffer containing the full length identifier (not guaranteed null terminated) takes priority over MCAUserIdentifier.

LongMCAUserIdLength

Length of string pointed to by LongMCAUserIdPtr - must be set if LongMCAUserIdPtr is set.

Remote User Identifier

Only applies to CLNTCONN/SVRCONN channel pairs. If no CLNTCONN Security Exit is defined then these three fields are initialized by the client MCA, so they might contain a user identifier from the environment of the client which can be used by a SVRCONN Security Exit for authentication and when specifying the MCA User Identifier. If a CLNTCONN Security Exit is defined then these fields are not initialized and can be set by the CLNTCONN Security Exit, or security messages can be used to pass a user identifier from Client to Server.

RemoteUserIdentifier

First 12 Bytes of identifier padded to the right with blanks.

LongRemoteUserIdPtr

Pointer to a buffer containing the full length identifier (not guaranteed null terminated) takes priority over RemoteUserIdentifier.

LongRemoteUserIdLength

Length of string pointed to by LongRemoteUserIdPtr - must be set if LongRemoteUserIdPtr is set.

PMQLONG pDataLength

input/output

Pointer to MQLONG. Contains the length of any Security Exit contained in the AgentBuffer upon invocation of the Security Exit. Must be set by a Security Exit to the length of any message being sent in the AgentBuffer or ExitBuffer.

PMQLONG pAgentBufferLength

input

Pointer to MQLONG. The length of the data contained in the AgentBuffer on invocation of the Security Exit.

PMQVOID pAgentBuffer

input/output

On invocation of the Security Exit, this points to any message sent from the partner exit. If ExitResponse2 in the MQCXP structure has the MQXR2_USE_AGENT_BUFFER flag set (default) then a Security Exit needs to set this parameter to point to any message data being sent.

PMQLONG pExitBufferLength

input/output

Pointer to MQLONG. This parameter is initialized to 0 on the first invocation of a Security Exit and the value returned is maintained between calls to the Security Exit during a security exchange.

PMQPTR pExitBufferAddr

input/output

This parameter is initialized to a null pointer on the first invocation of a Security Exit and the value returned is maintained between calls to the Security Exit during a security exchange. If the MQXR2_USE_EXIT_BUFFER flag is set in the ExitResponse2 in the MQCXP structure then a Security Exit needs to set this parameter to point to any message data being sent.

Differences in behavior between security exits defined on CLNTCONN/SVRCONN channel pairs and other channel pairs:

Security exits can be defined on all types of channel. However, the behavior of security exits defined on CLNTCONN/SVRCONN channel pairs is slightly different from security exits defined on other channel pairs.

A Security Exit on a CLNTCONN channel can set the Remote User Identifier in the channel definition for processing by a partner SVRCONN exit, or for OAM authorization if no SVRCONN Security Exit is defined and the MCAUSER field of the SVRCONN is not set.

If no CLNTCONN Security Exit is defined then the Remote User Identifier in the channel definition is set to a user identifier from the client environment (which can be blank) by the client MCA.

A security exchange between Security Exits defined on a CLNTCONN and SVRCONN channel pair completes successfully when the SVRCONN Security Exit returns an ExitResponse of MQXCC_OK. A security exchange between other channel pairs completes successfully when the Security Exit which initiated the exchange returns an ExitResponse of MQXCC_OK.

However, the MQXCC_SEND_AND_REQUEST_SEC_MSG ExitResponse code can be used to force continuation of the security exchange: If an ExitResponse of MQXCC_SEND_AND_REQUEST_SEC_MSG is returned by a CLNTCONN or SVRCONN Security Exit then the partner exit must respond by sending a security message (not MQXCC_OK or a null response) or the channel terminates. For Security Exits defined on other types of channel, an ExitResponse of MQXCC_OK returned in response to a MQXCC_SEND_AND_REQUEST_SEC_MSG from the partner Security Exit results in continuation of the security exchange as if a null response was returned and not in termination of the channel.

SSPI security exit:

IBM MQ for Windows supplies a security exit that provides authentication for IBM MQ channels by using the Security Services Programming Interface (SSPI). The SSPI provides the integrated security facilities of Windows.

This security exit is for both the IBM MQ client and the IBM MQ server.

The security packages are loaded from either security.dll or secur32.dll. These DLLs are supplied with your operating system.

One-way authentication is provided on Windows, using NTLM authentication services. Two-way authentication is provided on Windows 2000, using Kerberos authentication services.

The security exit program is supplied in source and object format. You can use the object code as it is, or you can use the source code as a starting point to create your own user-exit programs. For more information about using the object or source code of the SSPI security exit, see "Using the SSPI security

exit on Windows” on page 538

Channel send and receive exit programs:

You can use the send and receive exits to perform tasks such as data compression and decompression. You can specify a list of send and receive exit programs to be run in succession.

Channel send and receive exit programs are called at the following places in the processing cycle of an MCA:

- The send and receive exit programs are called for initialization at MCA initiation and for termination at MCA termination.
- The send exit program is invoked at one or other end of the channel, depending on the end at which a transmission for one message transfer is sent, immediately before a transmission is sent over the link. Note 4 explains why exits are available in both directions even though message channels send messages in one direction only.
- The receive exit program is invoked at one or other end of the channel, depending on the end at which a transmission for one message transfer is received, immediately after a transmission has been taken from the link. Note 4 explains why exits are available in both directions even though message channels send messages in one direction only.

There might be many transmissions for one message transfer, and there could be many iterations of the send and receive exit programs before a message reaches the message exit at the receiving end.

The channel send and receive exit programs are passed an agent buffer containing the transmission data as sent or received from the communications link. For send exit programs, the first 8 bytes of the buffer are reserved for use by the MCA, and must not be changed. If the program returns a different buffer, then these first 8 bytes must exist in the new buffer. The format of data presented to the exit programs is not defined.

A good response code must be returned by send and receive exit programs. Any other response causes an MCA abnormal end (abend).

Note: Do not issue an MQGET, MQPUT, or MQPUT1 call within sync point from a send or receive exit.

Note:

1. Send and receive exits typically work in pairs. For example a send exit might compress the data and a receive exit decompress it, or a send exit might encrypt the data and a receive exit decrypt it. When you define the appropriate channels, make sure that compatible exit programs are named for both ends of the channel.
2. If compression is turned on for the channel, the exits are passed compressed data.
3. Channel send and receive exits might be called for message segments other than for application data, for example, status messages. They are not called during the startup dialog, nor the security check phase.
4. Although message channels send messages in one direction only, channel-control data, such as heart beats and end of batch processing, flows in both directions, and these exits are available in both directions, also. However, some of the initial channel startup data flows are exempt from processing by any of the exits.
5. There are circumstances in which send and receive exits could be invoked out of sequence; for example, if you are running a series of exit programs or if you are also running security exits. Then, when the receive exit is first called upon to process data, it might receive data that has not passed through the corresponding send exit. If the receive exit just performed the operation, for example decompression, without first checking that it was required, the results would be unexpected.

You need to code your send and receive exits in such a way that the receive exit can check that the data it is receiving has been processed by the corresponding send exit. The recommended way to do so is to code your exit programs so that:

- The send exit sets the value of the ninth byte of data to 0 and shifts all the data along 1 byte, before performing the operation. (The first 8 bytes are reserved for use by the MCA.)
- If the receive exit receives data that has a 0 in byte 9, it knows that the data has come from the send exit. It removes the 0, performs the complementary operation, and shifts the resulting data back by 1 byte.
- If the receive exit receives data that has something other than 0 in byte 9, it assumes that the send exit has not run, and sends the data back to the caller unchanged.

When using security exits, if the channel is ended by the security exit it is possible that a send exit might be called without the corresponding receive exit. One way to prevent this problem is to code the security exit to set a flag, in `MQCD.SecurityUserData` or `MQCD.SendUserData`, for example, when the exit decides to end the channel. Then the send exit needs to check this field, and process the data only if the flag is not set. This check prevents the send exit from unnecessarily altering the data, and thus prevents any conversion errors that could occur if the security exit received altered data.

Channel send exit programs - reserving space:

You can use send and receive exits to transform the data before transmission. Channel send exit programs can add their own data about the transformation by reserving space in the transmission buffer.

This data is processed by the receive exit program and then removed from the buffer. For example, you might want to encrypt the data and add a security key for decryption.

How you reserve space and use it

When the send exit program is called for initialization, set the *ExitSpace* field of `MQXCP` to the number of bytes to be reserved. See `MQCXP` for details. *ExitSpace* can be set only during initialization, that is when *ExitReason* has the value `MQXR_INIT`. When the send exit is invoked immediately before transmission, with *ExitReason* set to `MQXR_XMIT`, *ExitSpace* bytes are reserved in the transmission buffer. *ExitSpace* is not supported on z/OS.

The send exit need not use all the reserved space. It can use less than *ExitSpace* bytes or, if the transmission buffer is not full, the exit can use more than the amount reserved. When setting the value of *ExitSpace*, you must leave at least 1 KB for message data in the transmission buffer. Channel performance can be affected if reserved space is used for large amounts of data.

The transmission buffer is normally 32Kb bytes long. However, if the channel uses TLS then the transmission buffer size is reduced to 15,352 bytes in order fit within the maximum record length defined by RFC 6101 and the related family of TLS standards. A further 1024 bytes are reserved for use by IBM MQ, so the maximum transmission buffer space usable by send exits is 14,328 bytes.

What happens at the receiving end of the channel

Channel receive exit programs must be set up to be compatible with the corresponding send exits. Receive exits must know the number of bytes in the reserved space and must remove the data in that space.

Multiple send exits

You can specify a list of send and receive exit programs to be run in succession. IBM MQ maintains a total for the space reserved by all the send exits. This total space must leave at least 1 KB for message data in the transmission buffer.

The following example shows how space is allocated for three send exits, called in succession:

1. When called for initialization:
 - Send exit A reserves 1 KB.
 - Send exit B reserves 2 KB.
 - Send exit C reserves 3 KB.
2. The maximum transmission size is 32 KB and the user data is 5 KB long.
3. Exit A is called with 5 KB of data; up to 27 KB are available, because 5 KB is reserved for exits B and C. Exit A adds 1 KB, the amount it reserved.
4. Exit B is called with 6 KB of data; up to 29 KB are available, because 3 KB is reserved for exit C. Exit B adds 1 KB, less than the 2 KB it reserved.
5. Exit C is called with 7 KB of data; up to 32 KB are available. Exit C adds 10K, more than the 3 KB it reserved. This amount is valid, because the total amount of data, 17 KB, is less than the 32 KB maximum.

The maximum transmission buffer size for a channel using TLS is 15,352 bytes, not 32Kb. This is because the underlying secure socket transmission segments are limited to 16Kb and some of the space is required for TLS record overheads. A further 1024 bytes are reserved for use by IBM MQ, so the maximum transmission buffer space usable by send exits is 14,328 bytes.

Channel message exit programs:

You can use the channel message exit to perform tasks such as encryption on the link, validation or substitution of incoming user IDs, message data conversion, journaling, and reference message handling. You can specify a list of message exit programs to be run in succession.

Channel message exit programs are called at the following places in the processing cycle of the MCA:

- At MCA initiation and termination
- Immediately after a sending MCA has issued an MQGET call
- Before the receiving MCA issues an MQPUT call



The message exit is passed an agent buffer containing the transmission queue header MQXQH, and the application message text as retrieved from the queue. The format of MQXQH is given in MQXQH - Transmission-queue header.

If you use reference messages (that is, messages that contain only a header that points to some other object that is to be sent), the message exit recognizes the header, MQRMH. It identifies the object, retrieves it in whatever way is appropriate appends it to the header, and passes it to the MCA for transmission to the receiving MCA. At the receiving MCA, another message exit recognizes that this message is a reference message, extracts the object, and passes the header on to the destination queue. See "Reference messages" on page 157 and "Running the Reference Message samples" on page 506 for more information about reference messages and some sample message exits that handle them.

Message exits can return the following responses:

- Send the message (GET exit). The message might have been changed by the exit. (This returns MQXCC_OK.)
- Put the message on the queue (PUT exit). The message might have been changed by the exit. (This returns MQXCC_OK.)
- Do not process the message. The message is placed on the dead-letter queue (undelivered message queue) by the MCA.
- Close the channel.
- Bad return code, which causes the MCA to abnormally end.

Note:

1. Message exits are called once for every complete message transferred, even when the message is split into parts.
2.   If you provide a message exit on UNIX or Linux, the automatic conversion of user IDs to lowercase characters (described here) does not operate.
3. An exit runs in the same thread as the MCA itself. It also runs inside the same unit of work (UOW) as the MCA because it uses the same connection handle. So any calls made under sync point are committed or backed out by the channel at the end of the batch. For example, one channel message exit program can send notification messages to another and these messages are only committed to the queue when the batch containing the original message is committed.
Therefore you can issue sync point MQI calls from a channel message exit program.

Message conversion outside the message exit:

Before calling the message exit, the receiving MCA performs some conversions on the message. This topic describes the algorithms used to perform the conversions.

Which headers are processed

A conversion routine runs in the MCA of the receiver before the message exit is called. The conversion routine begins with the MQXQH header at the beginning of the message. The conversion routine then processes through the chained headers that follow the MQXQH, performing conversion where necessary. The chained headers can extend beyond the offset contained in the HeaderLength parameter of the MQCXP data that is passed to the message exit of the receiver. The following headers are converted in-place:

- MQXQH (format name " MQXMIT ")
- MQMD (this header is part of the MQXQH and has no format name)
- MQMDE (format name " MQHMDE ")
- MQDH (format name " MQHDIST ")
- MQWIH (format name " MQHWIH ")

The following headers are not converted, but are stepped over as the MCA continues to process the chained headers:

- MQDLH (format name " MQDEAD ")
- any headers with format names beginning with the three characters 'MQH' (for example " MQHRF ") that are not otherwise mentioned

How the headers are processed

The Format parameter of each IBM MQ header is read by the MCA. The Format parameter is 8 bytes within the header, which are 8 single-byte characters containing a name.

The MCA then interprets the data following each header as being of the named type. If the Format is the name of a header type eligible for IBM MQ data conversion, it is converted. If it is another name indicating non-MQ data (for example MQFMT_NONE or MQFMT_STRING) then the MCA stops processing the headers.

What is the MQCXP HeaderLength?

The HeaderLength parameter in the MQCXP data supplied to a message exit is the total length of the MQXQH (which includes the MQMD), MQMDE and MQDH headers at the start of the message. These headers are chained using the 'Format' names and lengths.

MQWIH

Chained headers can extend beyond the HeaderLength into the user data area. The MQWIH header, if it is present, is one of those headers that appear beyond the HeaderLength.

If there is an MQWIH header in the chained headers, it is converted in-place before the message exit of the receiver is called.

Channel message retry exit program:

The channel message-retry exit is called when an attempt to open the target queue is unsuccessful. You can use the exit to determine under which circumstances to retry, how many times to retry, and how frequently.

This exit is also called at the receiving end of the channel at MCA initiation and termination.

The channel message-retry exit is passed an agent buffer containing the transmission queue header, MQXQH, and the application message text as retrieved from the queue. The format of MQXQH is given in Overview for MQXQH.

The exit is invoked for all reason codes; the exit determines for which reason codes it wants the MCA to retry, for how many times, and at what intervals. (The value of the message-retry count set when the channel was defined is passed to the exit in the MQCD, but the exit can ignore this value.)

The MsgRetryCount field in MQCXP is incremented by the MCA each time the exit is invoked, and the exit returns either MQXCC_OK with the wait time contained in the MsgRetryInterval field of MQCXP, or MQXCC_SUPPRESS_FUNCTION. Retries continue indefinitely until the exit returns MQXCC_SUPPRESS_FUNCTION in the ExitResponse field of MQCXP. See MQCXP for information about the action taken by the MCA for these completion codes.

If all the retries are unsuccessful, the message is written to the dead-letter queue. If there is no dead-letter queue available, the channel stops.

If you do not define a message-retry exit for a channel and a failure occurs that is likely to be temporary, for example MQRC_Q_FULL, the MCA uses the message-retry count and message-retry intervals set when the channel was defined. If the failure is of a more permanent nature and you have not defined an exit program to handle it, the message is written to the dead-letter queue.

Channel auto-definition exit program:

The channel auto-definition exit can be used when a request is received to start a receiver or server-connection channel but no definition for that channel exists (not for IBM MQ for z/OS). It can also be called on all platforms for cluster-sender and cluster-receiver channels to allow definition modification for an instance of the channel.

The channel auto-definition exit can be called on all platforms except z/OS when a request is received to start a receiver or server-connection channel but no channel definition exists. You can use it to modify the supplied default definition for an automatically defined receiver or server-connection channel, SYSTEM.AUTO.RECEIVER, or SYSTEM.AUTO.SVRCON. See Preparing channels for a description of how channel definitions can be created automatically.

The channel auto-definition exit can also be called when a request is received to start a cluster-sender channel. It can be called for cluster-sender and cluster-receiver channels to allow definition modification for this instance of the channel. In this case, the exit also applies to IBM MQ for z/OS. A common use of the channel auto-definition exit is to change the names of message exits (MSGEXIT, RCVEXIT, SCYEXIT, and SENDEXIT) because exit names have different formats on different platforms. If no channel

auto-definition exit is specified, the default behavior on z/OS is to examine a distributed exit name of the form *[path]/libraryname(function)* and take up to eight chars of function, if present, or libraryname. On z/OS, a channel auto-definition exit program must alter the fields addressed by `MsgExitPtr`, `MsgUserDataPtr`, `SendExitPtr`, `SendUserDataPtr`, `ReceiveExitPtr`, and `ReceiveUserDataPtr`, rather than the `MsgExit`, `MsgUserData`, `SendExit`, `SendUserData`, `ReceiveExit` and `ReceiveUserData` fields themselves.

For more information, see *Working with auto-defined channels*.

As with other channel exits, the parameter list is:

```
MQ_CHANNEL_AUTO_DEF_EXIT (ChannelExitParms, ChannelDefinition)
```

`ChannelExitParms` are described in MQCXP. `ChannelDefinition` is described in MQCD.

MQCD contains the values that are used in the default channel definition if they are not altered by the exit. The exit can modify only a subset of the fields; see `MQ_CHANNEL_AUTO_DEF_EXIT`. However, attempting to change other fields does not cause an error.

The channel auto-definition exit returns a response of either `MQXCC_OK` or `MQXCC_SUPPRESS_FUNCTION`. If neither of these responses is returned, the MCA continues processing as though `MQXCC_SUPPRESS_FUNCTION` were returned. That is, the auto-definition is abandoned, no new channel definition is created, and the channel cannot start.

Compiling channel exit programs on Windows, UNIX and Linux systems:

Use the following examples to help you compile channel-exit programs for Windows, UNIX and Linux systems.

Windows

Windows

The compiler and linker command for channel-exit programs on Windows:

```
cl.exe /Ic:\mqm\tools\c\include /nologo /c myexit.c
link.exe /nologo /dll myexit.obj /def:myexit.def /out:myexit.dll
```

UNIX Linux

UNIX and Linux systems

In these examples `exit` is the library name and `ChannelExit` is the function name. On AIX the export file is called `exit.exp`. These names are used by the channel definition to reference the exit program using the format described in MQCD- channel definition. See also the `MSGEXIT` parameter of the `DEFINE CHANNEL` command.

Sample compiler and linker commands for channel exits on AIX:

```
$ xlc_r -q64 -e MQStart -bE:exit.exp -bM:SRE -o /var/mqm/exits64/exit
exit.c -I/usr/mqm/inc
```

Sample compiler and linker commands for channel exits on HP-UX

```
$ c89 +DD64 +z -c -D HPUX_SOURCE -o exit.o exit.c -I/opt/mqm/inc
$ ld -b exit.o +ee MQStart +ee ChannelExit -o
/var/mqm/exits64/exit -L/usr/lib/pa20_64 -lpthread
$ rm exit.o
```

Sample compiler and linker commands for channel-exits on Linux where the queue manager is 32 bit:

```
$ gcc -shared -fPIC -o /var/mqm/exits/exit exit.c -I/opt/mqm/inc
```

Sample compiler and linker commands for channel-exits on Linux where the queue manager is 64-bit:

```
$ gcc -m64 -shared -fPIC -o /var/mqm/exits64/exit exit.c -I/opt/mqm/inc
```

Sample compiler and linker commands for channel exits on Solaris:

```
$ cc -xarch=v9 -mt -G -o /var/mqm/exits64/exit exit.c -I/opt/mqm/inc  
-R/usr/lib/64 -lsocket -lnsl -ldl
```

On the client, a 32 bit or 64 bit exit can be used. This exit must be linked to mqic_r.

On AIX, all functions that are called by IBM MQ must be exported. A sample export file for this make file:

```
#  
!channelExit  
MQStart
```

Configuring channel exits:

To call the channel exit, you must name it in the channel definition.

Channel exits must be named in the channel definition. You can do this naming when you first define the channels, or you can add the information later using, for example, the MQSC command ALTER CHANNEL. You can also give the channel exit names in the MQCD channel data structure. The format of the exit name depends on your IBM MQ platform; see MQCD or Script (MQSC) Commands for information.

If the channel definition does not contain a user-exit program name, the user exit is not called.

The channel auto-definition exit is the property of the queue manager, not the individual channel. In order for this exit to be called, it must be named in the queue manager definition. To alter a queue manager definition, use the MQSC command ALTER QMGR.

Writing data-conversion exits

This collection of topics contains information about how to write data-conversion exits.

Note: Not supported in MQSeries for VSE/ESA.

When you do an MQPUT, your application creates the message descriptor (MQMD) of the message. Because IBM MQ needs to be able to understand the contents of the MQMD regardless of the platform it is created on, it is converted automatically by the system.

Application data, however, is not converted automatically. If character data is being exchanged between platforms where the *CodedCharSetId* and *Encoding* fields differ, for example, between ASCII and EBCDIC, the application must arrange for conversion of the message. Application data conversion can be performed by the queue manager itself or by a user exit program, referred to as a *data-conversion exit*. The queue manager can perform data conversion itself, using one of its built-in conversion routines, if the application data is in one of the built-in formats (such as MQFMT_STRING). This topic contains information about the data-conversion exit facility that IBM MQ provides for when the application data is not in a built-in format.

Control can be passed to the data-conversion exit during an MQGET call. This avoids converting across different platforms before reaching the final destination. However, if the final destination is a platform that does not support data conversion on the MQGET, you must specify CONVERT(YES) on the sender channel that sends the data to its final destination. This ensures that IBM MQ converts the data during transmission. In this case, your data-conversion exit must reside on the system where the sender channel is defined.

The MQGET call is issued directly by the application. Set the *CodedCharSetId* and *Encoding* fields in the MQMD to the character set and encoding required. If your application uses the same character set and encoding as the queue manager, set *CodedCharSetId* to MQCCSI_Q_MGR, and *Encoding* to MQENC_NATIVE. After the MQGET call completes, these fields have the values appropriate to the message data returned. These might differ from the values required if the conversion was not successful. Your application should reset these fields to the values required before each MQGET call.


The conditions required for the data-conversion exit to be called are defined for the MQGET call in MQGET.

For a description of the parameters that are passed to the data-conversion exit, and detailed usage notes, see Data conversion for the MQ_DATA_CONV_EXIT call and the MQDXP structure.

Programs that convert application data between different machine encodings and CCSIDs must conform to the IBM MQ data conversion interface (DCI).

With the introduction of Multicast clients, API exits and data-conversion exits need to be able to run on the client-side because some messages might not go through the queue manager. The following libraries are now part of the client packages as well as the server packages:

Table 44. Libraries that are now in the client and server packages

Operating system	Libraries
Windows	32 bit & 64 bit: mqm.dll & mqm.pdb
Linux & HP-UX	32 bit & 64 bit: libmqm.so & libmqm_r.so
AIX	32 bit & 64 bit: libmqm.a & libmqm_r.a
Solaris	32 bit & 64 bit: libmqm.so
 IBM i	LIBMQM & LIBMQM_R

Invoking the data-conversion exit:

A data-conversion exit is a user-written exit that receives control during the processing of an MQGET call.

The exit is invoked if the following statements are true:

- The MQGMO_CONVERT option is specified on the MQGET call.
- Some or all of the message data is not in the requested character set or encoding.
- The *Format* field in the MQMD structure associated with the message is not MQFMT_NONE.
- The *BufferLength* specified on the MQGET call is not zero.
- The message data length is not zero.
- The message contains data that has a user-defined format. The user-defined format can occupy the entire message, or be preceded by one or more built-in formats. For example, the user-defined format might be preceded by an MQFMT_DEAD_LETTER_HEADER format. The exit is invoked to convert only the user-defined format; the queue manager converts any built-in formats that precede the user-defined format.

A user-written exit can also be invoked to convert a built-in format, but this happens only if the built-in conversion routines cannot convert the built-in format successfully.

There are some other conditions, described fully in the usage notes of the MQ_DATA_CONV_EXIT call in MQ_DATA_CONV_EXIT.

See MQGET for details of the MQGET call. Data-conversion exits cannot use MQI calls, other than MQXCNVC.

A new copy of the exit is loaded when an application attempts to retrieve the first message that uses that *Format* since the application connected to the queue manager. A new copy might also be loaded at other times if the queue manager has discarded a previously loaded copy.

The data-conversion exit runs in an environment like that of the program that issued the MQGET call. As well as user applications, the program can be an MCA (message channel agent) sending messages to a destination queue manager that does not support message conversion. The environment includes address space and user profile, where applicable. The exit cannot compromise the integrity of the queue manager, because it does not run in the queue manager's environment.

▶ z/OS

Data conversion on z/OS

On z/OS, be aware of the following:

- Exit programs can be written in assembly language only.
- Exit programs must be reentrant, and capable of running anywhere in storage.
- Exit programs must restore the environment on exit to that at entry, and must free any storage obtained.
- Exit programs must not WAIT, or issue ESTAEs or SPIEs.
- Exit programs are typically invoked as if by z/OS LINK in:
 - Non-authorized problem program state
 - Primary address space control mode
 - Non cross-memory mode
 - Non access-register mode
 - 31 bit addressing mode
 - TCB-PRB mode
- When used by a CICS application, the exit is invoked by EXEC CICS LINK, and must conform to the CICS programming conventions. The parameters are passed by pointers (addresses) in the CICS communication area (COMMAREA).

Although not recommended, user exit programs can also use CICS API calls, with the following caution:

- Do not issue sync points, as the results could influence units of work declared by the MCA.
- Do not update any resources controlled by a resource manager other than IBM MQ for z/OS, including those controlled by CICS Transaction Server.

For channels with CONVERT=YES, the exit is loaded from the data set referenced by the CSQXLIB DD statement. MQ-supplied exits CSQCBDCI and CSQCBDCO for the IBM MQ CICS Bridge are in SCSQAUTH.

Writing a data-conversion exit program for IBM i:

Information about steps to consider when writing MQ data-conversion exit programs for IBM i.

Follow these steps:

1. Name your message format. The name must fit in the *Format* field of the MQMD. The *Format* name must not have leading embedded blanks, and trailing blanks are ignored. The object's name must have no more than eight non-blank characters, because the *Format* is only eight characters long. Remember to use this name each time that you send a message (our example uses the name *Format*).
2. Create a structure to represent your message. See Valid syntax for an example.
3. Run this structure through the CVTMQMMDTA command to create a code fragment for your data-conversion exit.

The functions generated by the CVTMQMMDTA command use macros that are shipped in the file QMQM/H(AMQSVMHHA). These macros are written assuming that all structures are packed; amended them if this is not the case.

4. Take a copy of the supplied skeleton source file, QMQMSAMP/QCSRC(AMQSVFC4) and rename it. (Our example uses the name EXIT_MOD.)
5. Find the following comment boxes in the source file and insert code as described:

- a. Toward the end of the source file, a comment box starts with:

```
/* Insert the functions produced by the data-conversion exit */
```

Here, insert the code fragment generated in step 3.

- b. Near the middle of the source file, a comment box starts with:

```
/* Insert calls to the code fragments to convert the format's */
```

This is followed by a commented-out call to the function ConverttagSTRUCT.

Change the name of the function to the name of the function that you added in step 5a. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.

- c. Near the beginning of the source file, a comment box starts with:

```
/* Insert the function prototypes for the functions produced by */
```

Here, insert the function prototype statements for the functions added in step 5a.

If the message contains character data, the generated code calls MQXCNVC; this can be resolved by binding the service program QMQM/LIBMQM.

6. Compile the source module, EXIT_MOD, as follows:

```
CRTCMOD MODULE(library/EXIT_MOD) +  
SRCFILE(QCSRC) +  
TERASPACE(*YES *TSIFC)
```

7. Create/link the program.

For nonthreaded applications, use the following:

```
CRTPGM PGM(library/Format) +  
MODULE(library/EXIT_MOD) +  
BNDSRVPGM(QMQM/LIBMQM) +  
ACTGRP(QMQM) +  
USRPRF(*USER)
```

In addition to creating the data-conversion exit for the basic environment, another is required in the threaded environment. This loadable object must be followed by `_R`. Use the `LIBMQM_R` library to resolve calls to the `MQXCNVC`. Both loadable objects are required for a threaded environment.

```

CRTPGM PGM(library/Format_R) +
MODULE(library/EXIT_MOD) +
BNDSRVPGM(QMQM/LIBMQM_R) +
ACTGRP(QMQM) +
USRPRF(*USER)

```

- Place the output in the library list for the IBM MQ job. It is recommended that, for production, data-conversion exit programs be stored in QSYS.

Note:

- If CVTMQMDTA uses packed structures, all IBM MQ applications must use the `_Packed` qualifier.
- Data-conversion exit programs must be reentrant.
- MQXCNVC is the only MQI call that can be issued from a data-conversion exit.
- Compile the exit program with the user profile compiler option set to `*USER`, so that the exit runs with the authority of the user.
- Teraspace memory enablement is required for all user exits with IBM MQ for IBM i ; specify `TERASPACE(*YES *TSIFC)` in the `CRTCMOD` and `CRTBNDC` commands.

Writing a data-conversion exit program for IBM MQ for z/OS:

Information about steps to consider when writing data-conversion exit programs for IBM MQ for z/OS.

Follow these steps:

- Take the supplied source skeleton CSQ4BAX9 (for non-CICS environments) or CSQ4CAX9 (for CICS) as your starting point.
- Run the CSQUCVX utility.
- Follow the instructions in the prolog of CSQ4BAX9 or CSQ4CAX9 to incorporate the routines generated by the CSQUCVX utility, in the order that the structures occur in the message that you want to convert.
- The utility assumes that the data structures are not packed, that the implied alignment of the data is honored, and that the structures start on a fullword boundary, with bytes being skipped as required (as between ID and VERSION in the example in Valid syntax). If the structures are packed, omit the CMQXCALA macros that are generated. Therefore, consider declaring your structures in such a way that all fields are named and no bytes are skipped; in the example in Valid syntax, add a field `"MQBYTE DUMMY;"` between ID and VERSION.
- The supplied exit returns an error if the input buffer is shorter than the message format to be converted. Although the exit converts as many complete fields as possible, the error causes an unconverted message to be returned to the application. If you want to allow short input buffers to be converted as far as possible, including partial fields, change the TRUNC= value on the CSQXCDEFA macro to YES: no error is returned, so the application receives a converted message. The application must handle the truncation.
- Add any other special processing code that you need.
- Rename the program to your data format name.
- Compile and link-edit your program like a batch application program (unless it is for use with CICS applications). The macros in the code generated by the utility are in the library, **thlqual.SCSQMACS**. If the message contains character data, the generated code calls MQXCNVC. If your exit uses this call, link-edit it with the exit stub program CSQASTUB. The stub is language-independent and environment-independent. Alternatively, you can load the stub dynamically using the dynamic call name CSQXCNVC. See "Dynamically calling the IBM MQ stub" on page 427 for more information. Place the link-edited module in your application load library, and in a data set that is referenced by the CSQXLIB DD statement of your task procedure started by your channel initiator.
- If the exit is for use by CICS applications, compile and link-edit it like a CICS application program, including CSQASTUB if required. Place it in your CICS application program library. Define the program to CICS in the typical way, specifying EXECKEY(CICS) in the definition.

Note: Although the LE/370 runtime libraries are needed for running the CSQUCVX utility (see step 2 on page 368), they are not needed for link-editing or running the data-conversion exit itself (see steps 8 on page 368 and 9 on page 368).

See “Writing IMS bridge applications” on page 63 for information about data conversion within the IBM MQ - IMS bridge.

Writing a data-conversion exit for IBM MQ on UNIX and Linux systems:

Information about steps to consider when writing data-conversion exit programs for IBM MQ on UNIX and Linux systems.

Follow these steps:

1. Name your message format. The name must fit in the *Format* field of the MQMD, and be in uppercase, for example, MYFORMAT. The *Format* name must not have leading blanks. Trailing blanks are ignored. The object's name must have no more than eight non-blank characters, because the *Format* is only eight characters long. Remember to use this name each time that you send a message. If the data conversion exit is used in a threaded environment, the loadable object must be followed by `_r` to indicate that it is a threaded version.
2. Create a structure to represent your message. See Valid syntax for an example.
3. Run this structure through the `crtmqcvx` command to create a code fragment for your data-conversion exit.
The functions generated by the `crtmqcvx` command use macros that assume that all structures are packed; amend them if this is not the case.
4. Copy the supplied skeleton source file, renaming it to the name of your message format that you set in step 1. The skeleton source file, and the copy, are read-only.
The skeleton source file is called `amqsvfc0.c`.
5. On IBM MQ for AIX, a skeleton export file called `amqsvfc.exp` is also supplied. Copy this file, renaming it to `MYFORMAT.EXP`.
6. The skeleton includes a sample header file, `amqsvmha.h`, in the directory `MQ_INSTALLATION_PATH/inc`, where `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.. Make sure that your include path points to this directory to pick up this file.
The `amqsvmha.h` file contains macros that are used by the code generated by the `crtmqcvx` command. If the structure to be converted contains character data, these macros call `MQXCNVC`.
7. Find the following comment boxes in the source file and insert code as described:
 - a. Toward the end of the source file, a comment box starts with:

```
/* Insert the functions produced by the data-conversion exit */
```

Here, insert the code fragment generated in step 3.
 - b. Near the middle of the source file, a comment box starts with:

```
/* Insert calls to the code fragments to convert the format's */
```

This is followed by a commented-out call to the function `ConverttagSTRUCT`.
Change the name of the function to the name of the function that you added in step 7a. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.
 - c. Near the beginning of the source file, a comment box starts with:

```
/* Insert the function prototypes for the functions produced by */
```

Here, insert the function prototype statements for the functions added in step 3.
8. Compile your exit as a shared library, using `MQStart` as the entry point. To do this, see “Compiling data-conversion exits on UNIX and Linux systems” on page 370.

9. Place the output in the exit directory. The default exit directory is /var/mqm/exits for 32 bit systems and /var/mqm/exits64, for 64 bit systems. You can change these directories in the qm.ini or mqclient.ini file. This path can be set for each queue manager and the exit is only looked for in that path or paths.

Note:

1. If crtmqcvx uses packed structures, all IBM MQ applications must be compiled in this way.
2. Data-conversion exit programs must be reentrant.
3. MQXCNVC is the only MQI call that can be issued from a data-conversion exit.

Compiling data-conversion exits on UNIX and Linux systems:

Examples of how to compile a data conversion exit on UNIX and Linux systems.

On all platforms, the entry point to the module is MQStart.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

AIX

Compile the exit source code by issuing one of the following commands:

32 bit applications

Non-threaded

```
cc -e MQStart -bE:MYFORMAT.exp -bM:SRE -o /var/mqm/exits/MYFORMAT \
  MYFORMAT.c -I MQ_INSTALLATION_PATH/inc
```

Threaded

```
xlc_r -e MQStart -bE:MYFORMAT.exp -bM:SRE -o /var/mqm/exits/MYFORMAT_r \
  MYFORMAT.c -I MQ_INSTALLATION_PATH/inc
```

64 bit applications

Non-threaded

```
cc -q64 -e MQStart -bE:MYFORMAT.exp -bM:SRE -o /var/mqm/exits64/MYFORMAT \
  MYFORMAT.c -I MQ_INSTALLATION_PATH/inc
```

Threaded

```
xlc_r -q64 -e MQStart -bE:MYFORMAT.exp -bM:SRE -o /var/mqm/exits64/MYFORMAT_r \
  MYFORMAT.c -I MQ_INSTALLATION_PATH/inc
```

HP-UX Itanium platform

Compile and link the exit source code by issuing one of the following sets of commands:

32-bit applications

Non-threaded

Compile the exit source code:

```
c89 +e +z -c -D_HPUX_SOURCE -o MYFORMAT.o MYFORMAT.c -I MQ_INSTALLATION_PATH/inc
```

Link the exit object:

```
ld +b: -b MYFORMAT.o +ee MQStart -o \
  /var/mqm/exits/MYFORMAT -L/usr/lib/hpux32
rm MYFORMAT.o
```

Threaded

Compile the exit source code:

```
c89 +e +z -c -D_HPUX_SOURCE -o MYFORMAT.o MYFORMAT.c -I MQ_INSTALLATION_PATH/inc
```

Link the exit object:

```
ld +b: -b MYFORMAT.o +ee MQStart -o \  
    /var/mqm/exits/MYFORMAT_r -L/usr/lib/hpux32 \  
    -lpthread  
rm MYFORMAT.o
```

64-bit applications

Non-threaded

Compile the exit source code:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o MYFORMAT.o MYFORMAT.c -I MQ_INSTALLATION_PATH/inc
```

Link the exit object:

```
ld -b MYFORMAT.o +ee MQStart \  
    -o /var/mqm/exits64/MYFORMAT \  
    -L/usr/lib/hpux64  
rm MYFORMAT.o
```

Threaded

Compile the exit source code:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o MYFORMAT.o MYFORMAT.c -I MQ_INSTALLATION_PATH/inc
```

Link the exit object:

```
ld -b MYFORMAT.o +ee MQStart \  
    -o /var/mqm/exits64/MYFORMAT_r \  
    -L/usr/lib/hpux64 -lpthread  
rm MYFORMAT.o
```

Linux

Compile the exit source code by issuing one of the following commands:

31 bit applications

Non-threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/MYFORMAT MYFORMAT.c \  
    -I MQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/MYFORMAT_r MYFORMAT.c  
    -I MQ_INSTALLATION_PATH/inc
```

32 bit applications

Non-threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/MYFORMAT MYFORMAT.c  
    -I MQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/MYFORMAT_r MYFORMAT.c  
    -I MQ_INSTALLATION_PATH/inc
```

64 bit applications

Non-threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/MYFORMAT MYFORMAT.c  
    -I MQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/MYFORMAT_r MYFORMAT.c
-I MQ_INSTALLATION_PATH/inc
```

Solaris

Compile the exit source code by issuing one of the following commands:

32-bit applications

SPARC platform

```
cc -xarch=v8plus -KPIC -mt -G -o /var/mqm/exits/MYFORMAT \
MYFORMAT.c -I MQ_INSTALLATION_PATH/inc -R/usr/lib/32 -lsocket -lnsl -ldl
```

x86-64 platform

```
cc -xarch=386 -KPIC -mt -G -o /var/mqm/exits/MYFORMAT \
MYFORMAT.c -I MQ_INSTALLATION_PATH/inc -R/usr/lib/32 -lsocket -lnsl -ldl
```

64-bit applications

SPARC platform

```
cc -xarch=v9 -KPIC -mt -G -o /var/mqm/exits64/MYFORMAT \
MYFORMAT.c -I MQ_INSTALLATION_PATH/inc -R/usr/lib/64 -lsocket -lnsl -ldl
```

x86-64 platform

```
cc -xarch=amd64 -KPIC -mt -G -o /var/mqm/exits64/MYFORMAT \
MYFORMAT.c -I MQ_INSTALLATION_PATH/inc -R/usr/lib/64 -lsocket -lnsl -ldl
```

Writing a data-conversion exit for IBM MQ for Windows:

Information about steps to consider when writing data-conversion exit programs for IBM MQ for Windows.

Follow these steps:

1. Name your message format. The name must fit in the *Format* field of the MQMD. The *Format* name must not have leading blanks. Trailing blanks are ignored. The object's name must have no more than eight non-blank characters, because the *Format* is only eight characters long.

A .DEF file called amqsvfcn.def is also supplied in the samples directory, *MQ_INSTALLATION_PATH*\Tools\C\Samples. *MQ_INSTALLATION_PATH* is the directory where IBM MQ is installed. Take a copy of this file and rename it, for example, to MYFORMAT.DEF. Make sure that the name of the DLL being created and the name specified in MYFORMAT.DEF are the same. Overwrite the name FORMAT1 in MYFORMAT.DEF with the new format name.

Remember to use this name each time that you send a message.

2. Create a structure to represent your message. See Valid syntax for an example.
3. Run this structure through the crtmqcvx command to create a code fragment for your data-conversion exit.

The functions generated by the CRTMQCVX command use macros that are written assuming that all structures are packed; amend them if this is not the case.

4. Copy the supplied skeleton source file, amqsvfc0.c, renaming it to the name of your message format that you set in step 1.

amqsvfc0.c is in *MQ_INSTALLATION_PATH*\Tools\C\Samples where *MQ_INSTALLATION_PATH* is the directory where IBM MQ is installed. (The default installation directory is C:\Program Files\IBM\MQ.)

The skeleton includes a sample header file amqsvmha.h in the *MQ_INSTALLATION_PATH*\Tools\C\include directory. Make sure that your include path points to this directory to pick up this file.

The amqsvmha.h file contains macros that are used by the code generated by the CRTMQCVX command. If the structure to be converted contains character data, these macros call MQXCNV.

5. Find the following comment boxes in the source file and insert code as described:

- a. Toward the end of the source file, a comment box starts with:
`/* Insert the functions produced by the data-conversion exit */`

Here, insert the code fragment generated in step 3 on page 372.

- b. Near the middle of the source file, a comment box starts with:
`/* Insert calls to the code fragments to convert the format's */`

This is followed by a commented-out call to the function `ConverttagSTRUCT`.

Change the name of the function to the name of the function that you added in step 5a. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.

- c. Near the beginning of the source file, a comment box starts with:
`/* Insert the function prototypes for the functions produced by */`

Here, insert the function prototype statements for the functions added in step 3 on page 372.

6. Create the following command file:

```
c1 -I MQ_INSTALLATION_PATH\Tools\C\Include -Tp \
MYFORMAT.C
```

```
MYFORMAT.DEF
```

where `MQ_INSTALLATION_PATH` is the directory where IBM MQ is installed.

7. Issue the command file to compile your exit as a DLL file.
8. Place the output in the exit subdirectory below the IBM MQ data directory. The default directory for installing your exits on 32 bit systems is `MQ_DATA_PATH\Exits` and for 64 bit systems is `MQ_DATA_PATH\Exits64`

The path used to look for the data-conversion exits is given in the registry. The registry folder is:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\WebSphere MQ\Installation\MQ_INSTALLATION_NAME\Configuration\ClientExitPa
```

and the registry key is: `ExitsDefaultPath`. This path can be set for each queue manager and the exit is only looked for in that path or paths.

Note:

1. If `CRTMQCVX` uses packed structures, all IBM MQ applications must be compiled in this way.
2. Data-conversion exit programs must be reentrant.
3. `MQXCNVC` is the only MQI call that can be issued from a data-conversion exit.

Exit and switch load files on Windows operating systems:

The IBM WebSphere MQ for Windows Version 7.5 queue manager processes are 32-bit. As a result, when using 64-bit applications, some types of exit and XA switch load files also need to have a 32-bit version available for use by the queue manager. If the 32-bit version of the exit or XA switch load file is required and is not available, then the relevant API call or command fails.

Two attributes are supported in the `qm.ini` file for `ExitPath`. These are `ExitsDefaultPath=MQ_INSTALLATION_PATH\exits` and `ExitsDefaultPath64=MQ_INSTALLATION_PATH\exits64`. `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed. Using these ensures that the appropriate library can be found. If an exit is used in an IBM MQ cluster, this also ensures that the appropriate library on a remote system can be found.

The following table lists the different types of Exit and Switch load files and notes whether 32-bit or 64-bit versions, or both, are required, according to whether 32-bit or 64-bit applications are being used:

File types	32-bit applications	64-bit applications
API-crossing exit	32-bit	32-bit and 64-bit
Data conversion exit	32-bit	64-bit
Server Channel exits (all types)	32-bit	32-bit
Client Channel exits (all types)	32-bit	64-bit
Installable service exit	32-bit	32-bit
Service trace module	32-bit	32-bit and 64-bit
Cluster WLM exit	32-bit	32-bit
Pub/Sub routing exit	32-bit	32-bit
Database switch load files	32-bit	32-bit and 64-bit
External Transaction Manager AX libraries	32-bit	64-bit

Referencing connection definitions using a pre-connect exit from a repository

IBM MQ MQI clients can be configured to look up a repository to obtain connection definitions using a pre-connect exit library.

Introduction

A client application can connect to a queue manager using client channel definition tables (CCDT). Generally, the CCDT file is located on a central network file server, and have clients referencing it. Since it is difficult to manage and administer various client applications referencing the CCDT file, a flexible approach is to store the client definitions in a global repository like an LDAP directory, a WebSphere Registry and Repository or any other repository. Storing the client connection definitions in a repository makes managing client connection definitions easier, and applications can access the correct and most current client connection definitions.

During the MQCONN/X call execution, the IBM MQ MQI client loads an application specified pre-connect exit library, and invokes an exit function to retrieve connection definitions. The retrieved connection definitions are then used to establish connection to a queue manager. The details of exit library and function to invoke are specified in the `mqlclient.ini` configuration file.

Syntax

```
void MQ_PRECONNECT_EXIT (pExitParms, pQMgrName, ppConnectOpts, pCompCode, pReason);
```

Parameters

pExitParms

Type: PMQNX input /output

The **PreConnection** exit parameter structure.

The structure is allocated and maintained by the caller of the exit.

pQMgrName

Type: PMQCHAR input/output

Name of the queue manager.

On input, this parameter is the filter string supplied to the MQCONN API call through the **QMgrName** parameter. This field might be blank, explicit, or contain certain wildcard characters. The field is changed by the exit. The parameter is NULL when the exit is called with MQXR_TERM.

ppConnectOpts

Type: ppConnectOpts input/output

Options that control the action of MQCONN.

This is a pointer to an MQCNO connection options structure that controls the action of the MQCONN API call. The parameter is NULL when the exit is called with MQXR_TERM. The MQI client always provides an MQCNO structure to the exit, even if it was not originally provided by the application. If an application provides an MQCNO structure, the client makes a duplicate to pass it to the exit where it is modified. The client retains the ownership of the MQCNO.

An MQCD referenced through the MQCNO takes precedence over any connection definition provided through the array. The client uses the MQCNO structure to connect to the queue manager and the others are ignored.

pCompCode

Type: PMQLONG input/output

Completion code.

Pointer to an MQLONG that receives the exits completion code. It must be one of the following values:

- MQCC_OK - Successful completion
- MQCC_WARNING - Warning (partial completion)
- MQCC_FAILED - Call failed

pReason

Type: PMQLONG input/output

Reason qualifying pCompCode.

Pointer to an MQLONG that receives the exit reason code. If the completion code is MQCC_OK, the only valid value is:

- MQRC_NONE - (0, x'000') No reason to report.

If the completion code is MQCC_FAILED or MQCC_WARNING, the exit function can set the reason code field to any valid MQRC_* value.

C Invocation

```
void MQ_PRECONNECT_EXIT (&ExitParms, &QMgrName, &pConnectOpts, &CompCode, &Reason);
```

Parameter

```

PMQNX  pExitParms    /*PreConnect exit parameter structure*/
PMQCHAR pQMgrName    /*Name of the queue manager*/
PPMQCNO ppConnectOpts/*Options controlling the action of MQCONN*/
PMQLONG pCompCode    /*Completion code*/
PMQLONG pReason      /*Reason qualifying pCompCode*/

```

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 = *user_data*

This attribute specifies user data that is passed to the preconnect exit. The data that is passed to the preconnect exit is specific to the implementation of the preconnect exit that you are using and what data it is expecting to be passed.

This attribute can be read by C and unmanaged .NET clients.

For example, this attribute could be used to specify the URL of the repository where connection definitions are stored, such as, 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 = *myMod*

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

Writing and compiling publish exits

You can configure a publish exit at the queue manager to change the contents of a published message before it is received by subscribers. You can also change the message header, or not deliver the message to a subscription.

Note: Publish exits are not supported on z/OS.

You can use the publish exit to inspect and alter messages delivered to subscribers:

- Examine the contents of a message published to each subscriber
- Modify the contents of a message published to each subscriber
- Alter the queue to which a message is put
- Stop the delivery of a message to a subscriber

Writing a publish exit

Use the steps in “Writing exits and installable services on UNIX, Linux and Windows” on page 311, to help you write and compile your exit.

The provider of the publish exit defines what the exit does. The exit, however, must conform to the rules defined in MQPSXP.

IBM MQ does not provide an implementation of the MQ_PUBLISH_EXIT entry point. It does provide a C language typedef declaration. Use the typedef to declare the parameters to a user-written exit correctly. The following example illustrates how to use the typedef declaration:

```
#include "cmqec.h"

MQ_PUBLISH_EXIT MyPublishExit;

void MQENTRY MyPublishExit( PMQPSXP pExitParms,
                           PMQPBC  pPubContext,
                           PMQSBC  pSubContext )
{
  /* C language statements to perform the function of the exit */
}
```

The publish exit runs within the queue manager process, as a result of the following operations:

- A Publish operation where a message is delivered to one or more subscribers
- A Subscribe operation where one or more retained messages are delivered
- A Subscription Request operation where one or more retained messages are delivered

If the publish exit is called for a connection, the first time that it is called an *ExitReason* code of MQXR_INIT is set. Before the connection disconnects after using a publish exit, the exit is called with an *ExitReason* code of MQXR_TERM.

If the publish exit is configured, but cannot be loaded when the queue manager is started, publish/subscribe message operations are inhibited for the queue manager. You must fix the problem or restart the queue manager before publish/subscribe messaging is re-enabled.

Each IBM MQ connection that requires the publish exit might fail to load or initialize the exit. If the exit fails to load or initialize, publish/subscribe operations that require the publish exit are disabled for that connection. The operations fail with the IBM MQ reason code MQRC_PUBLISH_EXIT_ERROR.

The context in which the publish exit is called is the connection by an application to the queue manager. A user data area is maintained by queue manager for each connection that is performing publish operations. The exit can retain information in the user data area for each connection.

A publish exit can use some MQI calls. It can only use those MQI calls that manipulate message properties. The calls are:

- MQBUFMH
- MQCRTMH
- MQDLTMH
- MQDLTMP
- MQMHBUF
- MQINQMP
- MQSETMP

If the publish exit changes the destination queue manager or queue name, no new authority check is carried out.

Compiling a publish exit

The publish exit is a dynamically loaded library; it can be thought of as a channel-exit. For information about compiling exits, see “Writing exits and installable services on UNIX, Linux and Windows” on page 311.

Sample publish exit

The sample exit program is called `amqspse0.c`. It writes a different message to a log file depending on whether the exit was called for initialize, publish, or terminate operations. It also demonstrates the use of the exit user area field to allocate and free storage appropriately.

Configuring publish exits:

You must define certain attributes to configure a publish exit.

On Windows and Linux you can use the IBM MQ explorer to define the attributes. The attributes are defined on the queue manager properties page, under Publish/Subscribe.

To configure the publish exit in the `qm.ini` file on UNIX and Linux systems, create a stanza called `PublishSubscribe`. The `PublishSubscribe` stanza has the following attributes:

PublishExitPath=`[path]` | *module_name*

Module name and path containing the publish exit code. The maximum length of this field is `MQ_EXIT_NAME_LENGTH`. The default is no publish exit.

PublishExitFunction= *function_name*

Name of the function entry point into the module that contains the publish exit code. The maximum length of this field is `MQ_EXIT_NAME_LENGTH`.



On IBM i, if a program is used, omit `PublishExitFunction`.

PublishExitData= *string*

If the queue manager is calling a publish exit, it passes an `MQPSXP` structure as input. The data specified using the **PublishExitData** attribute is provided in the *ExitData* field of the structure. The string can be up to `MQ_EXIT_DATA_LENGTH` characters in length. The default is 32 blank characters.

Writing and compiling cluster workload exits

Write a cluster workload exit program to customize the workload management of clusters. You might take the cost of using a channel at different times of day, or message content, into account when routing messages. These are factors that are not considered by the standard workload management algorithm.

In most cases the workload management algorithm is sufficient for your needs. However, so that you can provide your own user-exit program to tailor workload management, IBM MQ includes a user exit, the cluster workload exit.

You might have some particular information about your network or messages that you could use to influence workload balancing. You might know which are the high-capacity channels or the cheap network routes, or you might want to route messages depending upon their content. You could decide to write a cluster workload exit program, or use one supplied by a third party.

The cluster workload exit is called when accessing a cluster queue. It is called by MQOPEN, MQPUT1 and MQPUT.

The target queue manager selected at MQOPEN time is fixed if MQ00_BIND_ON_OPEN is specified. In this case the exit is run only once.

If the target queue manager is not fixed at MQOPEN time, the target queue manager is chosen at the time of the MQPUT call. If the target queue manager is not available, or fails while the message is still on the transmission queue, the exit is called again. A new target queue manager is selected. If the message channel fails while the message is being transferred, and the message is backed out, a new target queue manager is selected.

Multi On Multiplatforms, the queue manager loads the new cluster workload exit the next time the queue manager is started.

If the queue manager definition does not contain a cluster workload exit program name, the cluster workload exit is not called.

Various data are passed to a cluster workload exit in the exit parameter structure, MQWXP:

- The message definition structure, MQMD.
- The message length parameter.
- A copy of the message, or part of the message.

On non-z/OS platforms, if you use CLWLMODE=FAST, each operating system process loads its own copy of the exit. Different connections to the queue manager can cause different copies of the exit to be invoked. If the exit is run in the default safe mode, CLWLMODE=SAFE, a single copy of the exit runs in its own separate process.

Writing cluster workload exits

z/OS For information about writing cluster workload exits for z/OS, see “Cluster workload exit programming for IBM MQ for z/OS” on page 381.

Multi For Multiplatforms, cluster workload exits must not use MQI calls. In other respects, the rules for writing and compiling cluster workload exit programs are like the rules that apply to channel exit programs. Follow the steps in “Writing exits and installable services on UNIX, Linux and Windows” on page 311, and use the sample program, “Sample cluster workload exit” on page 380 to help write and compile your exit.

For more information about channel exits, see “Writing channel-exit programs” on page 344.

Configuring cluster workload exits

You name cluster workload exits in the queue manager definition by specifying the cluster workload exit attribute on the ALTER QMGR command. For example:

```
ALTER QMGR CLWLEXIT(myexit)
```

Related information:

Cluster workload exit call and data structures

Sample cluster workload exit:

IBM MQ includes a sample cluster workload exit program. You can copy the sample and use it as a basis for your own programs.

z/OS IBM MQ for z/OS

The sample cluster workload exit program is supplied in Assembler and in C. The Assembler version is called CSQ4BAF1 and can be found in the library th1qual.SCSQASMS. The C version is called CSQ4BCF1 and can be found in the library th1qual.SCSQC37S. th1qual is the target library high-level qualifier for IBM MQ data sets in your installation.

Multi IBM MQ for Multiplatforms

The sample cluster workload exit program is supplied in C and is called amqswlm0.c. It can be found in:

Table 45. Sample cluster workload exit program location for Multiplatforms

Platform	Filepath
AIX, HP-UX, Sun Solaris	<i>MQ_INSTALLATION_PATH</i> /samp
Windows	<i>MQ_INSTALLATION_PATH</i> \Tools\c\Samples
IBM i	The qmqm library

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

This sample exit routes all messages to a particular queue manager, unless that queue manager becomes unavailable. It reacts to the failure of the queue manager by routing messages to another queue manager.

Indicate which queue manager you want messages to be sent to. Supply the name of the cluster-receiver channel in the CLWLDATA attribute on the queue manager definition. For example:

```
ALTER QMGR CLWLDATA(' my-cluster-name. my-queue-manager ')
```

To enable the exit, supply its full path and name in the CLWLEXIT attribute:

UNIX Linux On UNIX and Linux:

```
ALTER QMGR CLWLEXIT(' path /amqswlm(c1wlFunction)')
```

Windows On Windows:

```
ALTER QMGR CLWLEXIT(' path \amqswlm(c1wlFunction)')
```

z/OS On z/OS:

```
ALTER QMGR CLWLEXIT(CSQ4BxF1)
```

where x is either 'A' or 'C', depending on the programming language of the version you are using.

IBM i On IBM i, use either of the following commands:

- Use the MQSC command:

```
ALTER QMGR CLWLEXIT('AMQSWLM      library      ')
```

Both the program name and the library name occupy 10 characters and are blank-padded to the right if necessary.

- Use the CL command:

```
CHGMQM MQMNAME( qmgrname ) CLWLEXIT(' library /AMQSWLM')
```

Now, instead of using the supplied workload management algorithm, IBM MQ calls this exit to route all messages to your chosen queue manager.

Cluster workload exit programming for IBM MQ for z/OS:

Cluster workload exits are invoked as if by a z/OS **LINK** command. Exits are subject to a number of stringent programming rules. Avoid using most SVC commands that involve waits, or using a STAE or ESTAE in a workload exit.

Cluster workload exits are invoked as if by a z/OS **LINK** in:

- Non-authorized problem program state
- Primary address space control mode
- Non-cross-memory mode
- Non-access register mode
- 31 bit addressing mode
- Storage key 8
- Program Key Mask 8
- TCB key 8

Put the link-edited modules in the data set specified by the CSQXLIB DD statement of the queue manager address space procedure. The names of the load modules are specified as the workload exit names in the queue manager definition.

When writing workload exits for IBM MQ for z/OS, the following rules apply:

- You must write exits in assembler or C. If you use C, it must conform to the C systems programming environment for system exits, described in the *z/OS C/C++ Programming Guide*, SC09-4765.
- If using the MQXCLWLN call, link edit with CSQMFLW, supplied in *thlqual.SCSQLLOAD*.
- Exits are loaded from the non-authorized libraries defined by a CSQXLIB DD statement. Providing CSQXLIB has DISP=SHR, exits can be updated while the queue manager is running, with the new version used in the next MQCONN thread the queue manager starts.
- Exits must be reentrant, and capable of running anywhere in virtual storage.
- Exits must reset the environment on return to that at entry.
- Exits must free any storage obtained, or ensure that storage is freed by a subsequent exit invocation.
- No MQI calls are allowed.
- Exits must not use any system services that could cause a wait, because a wait severely degrades the performance of the queue manager. In general, therefore, avoid an SVC, PC, or I/O.
- Exits must not issue an ESTAE or SPIE, apart from within any subtasks they attach.

Note: There are no absolute restrictions on what you can do in an exit. However, most SVCs involve waits, so avoid them, except for the following commands:

- **GETMAIN / FREEMAIN**
- **LOAD / DELETE**

Do not use ESTAEs and ESPIEs because their error handling might interfere with the error handling performed by IBM MQ. IBM MQ might not be able to recover from an error, or your exit program might not receive all the error information.

The system parameter EXITLIM limits the amount of time an exit might run for. The default value for EXITLIM is 30 seconds. If you see the return code MQRC_CLUSTER_EXIT_ERROR, 2266 X'8DA' your exit might be looping. If you think the exit needs more than 30 seconds to complete, increase the value of EXITLIM.

Building a procedural application

You can write an IBM MQ application in one of several procedural languages, and run the application on several different platforms.

Building your procedural application on AIX



The AIX publications describe how to build executable applications from the programs that you write.

This topic describes the additional tasks, and the changes to the standard tasks, that you must perform when building IBM MQ for AIX applications to run under AIX. C, C++, and COBOL are supported. For information about preparing your C++ programs, see Using C++.

The tasks that you must perform to create an executable application using IBM MQ for AIX vary with the programming language that your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the IBM MQ for AIX include files for the language that you are using. Make yourself familiar with the contents of these files. See “IBM MQ data definition files” on page 71 for a full description.

When you run threaded server or threaded client applications, set the environment variable AIXTHREAD_SCOPE=S.

Preparing C programs in AIX:

This topic contains information about linking libraries necessary to prepare C programs on AIX.

Precompiled C programs are supplied in the *MQ_INSTALLATION_PATH*/samp/bin directory. Use the ANSI compiler and run the following commands. For more information about programming 64-bit applications, see Coding standards on 64-bit platforms.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

For 32-bit applications:

```
$ xlc_r -o amqsput_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -lmqm
```

where amqsput0 is a sample program.

For 64-bit applications:

```
$ xlc_r -q64 -o amqsput_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -lmqm
```

where amqsput0 is a sample program.

If you are using the VisualAge C/C++ compiler for C++ programs, you must include the option `-q namemangling=v5` to get all the IBM MQ symbols resolved when linking the libraries.

If you want to use the programs on a machine that has only the IBM MQ MQI client for AIX installed, recompile the programs to link them with the client library (`-lmqic`) instead.

Linking libraries

You need the following libraries:

- Link your programs with the appropriate library provided by IBM MQ.

In a non-threaded environment, link to one of the following libraries:

Library file	Program/exit type
libmqm.a	Server for C
libmqic.a & libmqm.a	Client for C

In a threaded environment, link to one of the following libraries:

Library file	Program/exit type
libmqm_r.a	Server for C
libmqic_r.a & libmqm_r.a	Client for C

For example, to build a simple threaded IBM MQ application from a single compilation unit run the following commands.

For 32-bit applications:

```
$ xlc_r -o amqsputc_32_r amqsput0.c -I MQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -lmqm_r
```

where amqsput0 is a sample program.

For 64-bit applications:

```
$ xlc_r -q64 -o amqsputc_64_r amqsput0.c -I MQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -lmqm_r
```

where amqsput0 is a sample program.

If you want to use the programs on a machine that has only the IBM MQ MQI client for AIX installed, recompile the programs to link them with the client library (-lmqic) instead.

Note:

1. You cannot link to more than one library. That is, you cannot link to both a threaded and a non-threaded library at the same time.
2. If you are writing an installable service (see the Administering for further information), you need to link to the libmqmzf.a library in a non-threaded application and to the libmqmzf_r.a library in a threaded application.
3. If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries, Encina, or BEA Tuxedo, you need to link to the libmqmxa.a (or libmqmxa64.a if your transaction manager treats the 'long' type as 64 bit) and libmqz.a libraries in a non-threaded application and to the libmqmxa_r.a (or libmqmxa64_r.a) and libmqz_r.a libraries in a threaded application.
4. You need to link trusted applications to the threaded IBM MQ libraries. However, only one thread in a trusted application on IBM MQ on UNIX and Linux systems can be connected at a time.
5. You must link IBM MQ libraries before any other product libraries.

Preparing COBOL programs in AIX:

Use this information when preparing COBOL programs in AIX using IBM COBOL Set and Micro Focus COBOL.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

- 32 bit COBOL copy books are installed in the following directory:

MQ_INSTALLATION_PATH/inc/cobcpy32

and symbolic links are created in:

MQ_INSTALLATION_PATH/inc

- 64 bit COBOL copy books are installed in the following directory:

MQ_INSTALLATION_PATH/inc/cobcpy64

In the following examples set the **COBCPY** environment variable to:

MQ_INSTALLATION_PATH/inc/cobcpy32

for 32 bit applications, and:

MQ_INSTALLATION_PATH/inc/cobcpy64

for 64 bit applications.

You need to link your program with one of the following library files:

Library file	Program/exit type
libmqmcb.a	Server for COBOL (unthreaded application)
libmqmcb_r.a	Server for COBOL (threaded application)
libmqicb.a	Client for COBOL (unthreaded application)
libmqicb_r.a	Client for COBOL (threaded application)

You can use the IBM COBOL Set compiler or Micro Focus COBOL compiler depending on the program:

- Programs beginning amqm are suitable for the Micro Focus COBOL compiler, and
- Programs beginning amq0 are suitable for either compiler.

Preparing COBOL programs using IBM COBOL Set for AIX

Sample COBOL programs are supplied with IBM MQ. To compile such a program, enter the appropriate command from the following list:

32 bit non-threaded server application

```
$ cob2 -o amq0put0 amq0put0.cb1 -L MQ_INSTALLATION_PATH/lib -lmqmc -qLIB \  
-ICOBCPY_VALUE
```

32 bit non-threaded client application

```
$ cob2 -o amq0put0 amq0put0.cb1 -L MQ_INSTALLATION_PATH/lib -lmqic -qLIB \  
-ICOBCPY_VALUE
```

32 bit threaded server application

```
$ cob2_r -o amq0put0 amq0put0.cb1 -qTHREAD -L MQ_INSTALLATION_PATH/lib \  
-lmqmc_r -qLIB -ICOBCPY_VALUE
```

32 bit threaded client application

```
$ cob2_r -o amq0put0 amq0put0.cb1 -qTHREAD -L MQ_INSTALLATION_PATH/lib \  
-lmqic_r -qLIB -ICOBCPY_VALUE
```


64 bit non-threaded server application

```
$ cob2 -o amq0put0 amq0put0.cbl -q64 -L MQ_INSTALLATION_PATH/lib -lmqmcbl \
-qLIB -ICOBPCPY_VALUE
```

64 bit non-threaded client application

```
$ cob2 -o amq0put0 amq0put0.cbl -q64 -L MQ_INSTALLATION_PATH/lib -lmqicbl \
-qLIB -ICOBPCPY_VALUE
```

64 bit threaded server application

```
$ cob2_r -o amq0put0 amq0put0.cbl -q64 -qTHREAD -L MQ_INSTALLATION_PATH/lib \
-lmqmcb_r -qLIB -ICOBPCPY_VALUE
```

64 bit threaded client application

```
$ cob2_r -o amq0put0 amq0put0.cbl -q64 -qTHREAD -L MQ_INSTALLATION_PATH/lib \
-lmqicb_r -qLIB -ICOBPCPY_VALUE
```

Preparing COBOL programs using Micro Focus COBOL

Set environment variables before compiling your program as follows:

```
export COBPCPY=COBPCPY_VALUE
export LIBPATH=MQ_INSTALLATION_PATH/lib:$LIBPATH
```

To compile a 32 bit COBOL program using Micro Focus COBOL, enter:

- Server for COBOL

```
$ cob32 -xvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib -lmqmcbl
```

- Client for COBOL

```
$ cob32 -xvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib -lmqicbl
```

- Threaded Server for COBOL

```
$ cob32 -xtvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib -lmqmcb_r
```

- Threaded Client for COBOL

```
$ cob32 -xtvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb_r
```

To compile a 64 bit COBOL program using Micro Focus COBOL, enter:

- Server for COBOL

```
$ cob64 -xvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmcbl
```

- Client for COBOL

```
$ cob64 -xvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicbl
```

- Threaded Server for COBOL

```
$ cob64 -xtvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmcb_r
```

- Threaded Client for COBOL

```
$ cob64 -xtvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb_r
```

where amqminqx is a sample program

See the Micro Focus COBOL documentation for a description of the environment variables that you need to set up.

Preparing CICS application programs in AIX:

Use this information when preparing CICS programs in AIX.

Use *XA switch* modules to link CICS with IBM MQ. For more information on the XA switch structure, see The XA switch structures.

The sample source code file is provided to enable you to develop the XA switches for other transaction messages. The name of the switch load module provided is listed in Table 46.

Table 46. Essential code for CICS application programs on AIX: XA initialization routine

Description	C (source)	C (exec) - add to your XAD.Stanza
XA initialization routine	amqzscix.c	amqzsc - CICS for AIX

Use the prebuilt version of the IBM MQ switch load file *amqzsc*, which is provided with the product.

Always link your C transactions with the threadsafe IBM MQ library *libmqm_r.a.*, and your COBOL transactions with the COBOL library *libmqmcb_r.a.*

You can find more information about supporting CICS transactions in the Administering IBM MQ System Administration Guide.

TXSeries CICS support:

IBM MQ on AIX supports TXSeries CICS using the XA interface. Ensure that CICS applications are linked to the threaded version of the IBM MQ libraries.

You can run CICS programs using IBM COBOL Set for AIX or Micro Focus COBOL. The following sections describe the difference between running CICS programs on IBM COBOL Set for AIX and Micro Focus COBOL.

Write IBM MQ programs that are loaded into the same CICS region in either C or COBOL. You cannot make a combination of C and COBOL MQI calls into the same CICS region. Most MQI calls in the second language used fail with a reason code of `MQRC_HOBJ_ERROR`.

Preparing CICS COBOL programs using IBM COBOL Set for AIX

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

To use IBM COBOL, follow these steps:

1. Export the following environment variable:

```
export LDFlags="-qLIB -bI:/usr/lpp/cics/lib/cicsprIBMCOB.exp \  
-I MQ_INSTALLATION_PATH/inc -I/usr/lpp/cics/include \  
-e _iwz_cobol_main \  
"
```

where LIB is a compiler directive.

2. Translate, compile, and link the program by typing:

```
cicstcl -l IBMCOB yourprog.ccp
```

Preparing CICS COBOL programs using Micro Focus COBOL

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

To use Micro Focus COBOL, follow these steps:

1. Add the IBM MQ COBOL runtime library module to the runtime library using the following command:

```
cicsmkcobol -L/usr/lib/dce -L MQ_INSTALLATION_PATH/lib \  
            MQ_INSTALLATION_PATH/lib/libmqmcbt.o -lmqe_r
```

Note: With `cicsmkcobol`, IBM MQ does not allow you to make MQI calls in the C programming language from your COBOL application.

If your existing applications have any such calls, you are recommended to move these functions from the COBOL applications to your own library, for example, `myMQ.so`. After moving the functions, do not include the IBM MQ library `libmqmcbt.o` when building the COBOL application for CICS.

Additionally, if your COBOL application does not make any COBOL MQI call, do not link `libmqmz_r` with `cicsmkcobol`.

This creates the Micro Focus COBOL language method file and enables the CICS runtime COBOL library to call IBM MQ on UNIX and Linux systems.

Note: Run `cicsmkcobol` only when you install one of the following products:

- New version or release of Micro Focus COBOL
 - New version or release of CICS for AIX
 - New version or release of any supported database product (for COBOL transactions only)
 - New version or release of IBM MQ
2. Export the following environment variable:
`COBCPY= MQ_INSTALLATION_PATH/inc export COBCPY`
 3. Translate, compile, and link the program by typing:
`cicstcl -l COBOL -e yourprog.ccp`

Preparing CICS C programs

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Build CICS C programs using the standard CICS facilities:

1. Export *one* of the following environment variables:
 - `LD_FLAGS = "-L/ MQ_INSTALLATION_PATH lib -lmqm_r" export LD_FLAGS`
 - `USERLIB = "-L MQ_INSTALLATION_PATH lib -lmqm_r" export USERLIB`
2. Translate, compile, and link the program by typing:
`cicstcl -l C amqscic0.ccs`

CICS C sample transaction

Sample C source for a AIX IBM MQ transaction is provided by `AMQSCIC0.CCS`. The transaction reads messages from the transmission queue `SYSTEM.SAMPLE.CICS.WORKQUEUE` on the default queue manager and places them onto the local queue with a queue name that is contained in the transmission header of the message. Any failures are sent to the queue `SYSTEM.SAMPLE.CICS.DLQ`. Use the sample MQSC script `AMQSCIC0.TST` to create these queues and sample input queues.

Building your procedural application on HP Integrity NonStop Server

NSS Client

This information describes the additional tasks, and the changes to the standard tasks, that you must perform when you build IBM MQ client for HP Integrity NonStop Server applications to run under HP Integrity NonStop Server.

C, COBOL, and pTAL are supported.

OSS and Guardian headers and public libraries:

Provides lists of OSS and Guardian headers and public libraries. Listed are OSS headers, OSS public executable and public import libraries, Guardian headers, and Guardian public executable and public import libraries.

“OSS headers”

“OSS public executable and public import libraries”

“Guardian headers” on page 389

“Guardian public executable and public import libraries” on page 389

OSS headers

Table 47. OSS headers

Object	Location	Description
cmqbc.h	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ C-language header (OSS)
cmqc.h	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ C-language header (OSS)
cmqfc.h	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ C-language header (OSS)
cmqec.h	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ C-language header (OSS)
cmqpvc.h	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ C-language header (OSS)
cmqxc.h	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ C-language header (OSS)
cmqzc.h	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ C-language header (OSS)
cmqcobol.cpy	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ COBOL copybook (OSS)
cmqbt.tal	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ pTAL header (OSS)
cmqcft.tal	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ pTAL header (OSS)
cmqpst.tal	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ pTAL header (OSS)
cmqt.tal	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ pTAL header (OSS)
cmqxt.tal	<i>MQ_INSTALLATION_PATH/inc</i>	IBM MQ pTAL header (OSS)

OSS public executable and public import libraries

Table 48. OSS public executable and public import libraries

Object	Location	Description
libmqic.so	<i>MQ_INSTALLATION_PATH</i> /bin	IBM MQ public executable library (OSS unthreaded)
libmqic_r.so	<i>MQ_INSTALLATION_PATH</i> /bin	IBM MQ public executable library (OSS multi-threaded)
libmqic.so	<i>MQ_INSTALLATION_PATH</i> /lib	IBM MQ public import library (OSS unthreaded)
libmqic_r.so	<i>MQ_INSTALLATION_PATH</i> /lib	IBM MQ public import library (OSS multi-threaded)
mqicb	<i>MQ_INSTALLATION_PATH</i> /lib	IBM MQ public import library for COBOL (OSS)

Guardian headers

Table 49. Guardian headers

Object	Location	Description
cmqbcch	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ C-language header (Guardian)
cmqch	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ C-language header (Guardian)
cmqfch	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ C-language header (Guardian)
cmqech	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ C-language header (Guardian)
cmqpsch	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ C-language header (Guardian)
cmqxch	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ C-language header (Guardian)
cmqzch	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ C-language header (Guardian)
cmqcobol	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ COBOL copybook (Guardian)
cmqbt	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ pTAL header (Guardian)
cmqcft	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ pTAL header (Guardian)
cmqpst	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ pTAL header (Guardian)
cmqt	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ pTAL header (Guardian)
cmqxt	<i>MQ_INSTALLATION_PATH</i> /inc/G	IBM MQ pTAL header (Guardian)

Guardian public executable and public import libraries

Table 50. Guardian public executable and public import libraries

Object	Location	Description
mqic	<i>MQ_INSTALLATION_PATH</i> /bin/G	IBM MQ public executable library (Guardian)
mqicb	<i>MQ_INSTALLATION_PATH</i> /lib/G	IBM MQ public import library for COBOL (Guardian)

Preparing C programs in HP Integrity NonStop Server:

This topic contains information to consider when you are preparing C programs in HP Integrity NonStop Server together with examples of the commands you use when you are building applications when you are using the OSS C compiler and when you are using the Guardian C compiler.

Precompiled C programs are supplied in the *MQ_INSTALLATION_PATH*/opt/mqm/samp/bin directory. To build a sample from source code, use the c89 compiler.

You must link your programs with the appropriate library provided by IBM MQ. The following table lists the libraries that you must link to when you are preparing C programs on HP Integrity NonStop Server.

Table 51. . HP Integrity NonStop Server link libraries

Library	Description
libmqic.so	OSS unthreaded
libmqic_r.so	OSS multi-threaded
mqic	Guardian

Multi-threaded native IBM MQ applications must use the Posix User Threads (PUT) feature. There is no support for Standard Posix Threads (SPT) in this product.

Building applications using the OSS C compiler

This section contains examples of the commands that are used to build programs that are targeted for either OSS or Guardian when you are using the OSS compiler.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

The following example builds an unthreaded C client OSS application:

```
c89 -wsystype=oss -o amqsputc amqsput0.c -IMQ_INSTALLATION_PATH/opt/mqm/inc
-LMQ_INSTALLATION_PATH/opt/mqm/lib -lmqic
```

The following example builds a multi-threaded C client OSS application:

```
c89 -wsystype=oss -D_PUT_MODEL_ -o amqsputc amqsput0.c -IMQ_INSTALLATION_PATH/opt/mqm/inc
-LMQ_INSTALLATION_PATH/opt/mqm/lib -lmqic_r -lput
```

The following example builds a Guardian C client application:

```
c89 -wsystype=guardian -o /G/vol/subvol/amqsputc amqsput0.c -IMQ_INSTALLATION_PATH/opt/mqm/inc
-LMQ_INSTALLATION_PATH/opt/mqm/lib/G -lmqic
```

Building applications using the Guardian C compiler

This section contains examples of the commands that are used to build programs that are targeted for Guardian when you are using the Guardian compiler.

`MQ_INSTALLATION_PATH` represents the Guardian volume and subvolume in which IBM MQ is installed.

The following example builds a Guardian C client application:

```
CCOMP /in AMQSPUT0/ AMQSPUTC;&
    runnable,systype guardian,nolist,&
    ssv0 "$system.system",&
    ssv1 "MQINSTALLATION_SUBVOL",&
    ld(-LMQINSTALLATION_SUBVOL -lmqic)
```

Preparing COBOL programs for the IBM MQ client for HP Integrity NonStop Server:

This topic contains information to consider when you are preparing C programs for the IBM MQ client for HP Integrity NonStop Server. It contains examples of the commands you use when you are building applications when you are using the OSS ECOBOL compiler and when you are using the Guardian ECOBOL compiler.

To build a COBOL sample from source code, use the ECOBOL compiler.

The following table lists the libraries that are needed when you are preparing COBOL programs on HP Integrity NonStop Server . You must link your programs with the appropriate library provided by IBM MQ .

Table 52. . HP Integrity NonStop Server link libraries

Library	Description
libmqic.so	OSS unthreaded
mqic	Guardian

When you run a COBOL application that connects to a queue manager, you must first set the `SAVE-ENVIRONMENT` variable to ON. To set the `SAVE-ENVIRONMENT` variable to ON:

- For OSS, enter the following command:
`export SAVE-ENVIRONMENT=ON`
- For Guardian, enter the following command:
`param SAVE-ENVIRONMENT ON`

If you do not set the `SAVE-ENVIRONMENT` variable to ON, when the application attempts to connect to a queue manager, it fails with reason code 2058 (080A) (RC2058): `MQRC_Q_MGR_NAME_ERROR`.

Building applications using the OSS ECOBOL compiler

This section contains examples of the commands that are used to build programs that are targeted for either OSS or Guardian when you are using the OSS ECOBOL compiler.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

The following example builds a COBOL client OSS application:

```
ecobol -wsystype=oss
    -wcobol="ansi;port"
    -wcobol="consult MQ_INSTALLATION_PATH/opt/mqm/lib/mqicb"
    -wcopylib=MQ_INSTALLATION_PATH/opt/mqm/inc/cmqcobol.cpy
    -LMQ_INSTALLATION_PATH/opt/mqm/lib -lmqic
    -o amq0put0
    MQ_INSTALLATION_PATH/opt/mqm/samp/amq0put0.cb1
```

The following example builds a COBOL client Guardian application:

```
ecobol -wsystype=guardian
        -wcobol="ansi;port;save all"
        -wcobol="consult MQ_INSTALLATION_PATH/opt/mqm/lib/mqicb"
        -wcopylib=MQ_INSTALLATION_PATH/opt/mqm/inc/cmqcobol.cpy
        -LMQ_INSTALLATION_PATH/opt/mqm/lib/G -lmqic
        -o amq0put0
        MQ_INSTALLATION_PATH/opt/mqm/samp/amq0put0.cb1
```

Building applications using the Guardian ECOBOL compiler

This section contains examples of the commands that are used to build programs that are targeted for Guardian when you are using the Guardian ECOBOL compiler.

MQ_INSTALLATION_SUBVOL represents the Guardian volume and subvolume in which IBM MQ is installed.

The following example builds a COBOL client Guardian application:

```
ECOBOL /in MQSPUTL/ MQSPUT, MQINSTALLATION_SUBVOL.cmqcobol;
        call-shared;ansi;port;save all;nolist;runnable;
        consult MQINSTALLATION_SUBVOL.mqicb;
        eld(-LMQINSTALLATION_SUBVOL -lmqic)
```

Preparing pTAL programs:

Learn to build pTAL programs for the IBM MQ client on the HP Integrity NonStop Server platform.

To build a pTAL sample from source code, use the EPTAL compiler.

Note:

- pTAL IBM MQ applications must use a main routine that is written in either the C or COBOL languages.
- pTAL applications can be built only in Guardian.

The following table lists the library that is needed when you are preparing pTAL programs on HP Integrity NonStop Server. You must link your programs with the appropriate library provided by IBM MQ.

Table 53. . HP Integrity NonStop Server link library

Library	Description
mqic	Guardian

Building applications using the Guardian EPTAL compiler

This section contains examples of the commands that are used to build programs that are targeted for Guardian when you are using the Guardian EPTAL compiler.

MQINSTALLATION_SUBVOL represents the Guardian volume and subvolume in which IBM MQ is installed.

pTAL IBM MQ applications must use a main routine that is written in either the C or COBOL languages.

The following example builds a pTAL client Guardian application:

```
ASSIGN SSV0, $SYSTEM.SYSTEM
ASSIGN SSV1, MQINSTALLATION_SUBVOL

EPTAL /in MQINSTALLATION_SUBVOL.MQSPUTT/ MQSPUTO;nolist

CCOMP /in MQINSTALLATION_SUBVOL.MQSPBTC/ MQSPUT;
        runnable,systype guardian,extensions,nolist,
```



```

ssv0 "$system.system",
ssv1 "MQINSTALLATION_SUBVOL",
e1d(MQSPUTO -LMQINSTALLATION_SUBVOL -lmqic)

```

Building your procedural application on HP-UX

This information describes the additional tasks, and the changes to the standard tasks, that you must perform when building IBM MQ for HP-UX applications to run under HP-UX.

C, C++, and COBOL are supported. For information about preparing your C++ programs, see *Using C++*.

The tasks that you must perform to create an executable application using IBM MQ for HP-UX vary with the programming language that your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the IBM MQ for HP-UX include files for the language that you are using. Make yourself familiar with the contents of these files. See “IBM MQ data definition files” on page 71 for a full description.

Throughout this topic, we use the backslash (\) character to split long commands over more than one line. Do not enter this character; enter each command as a single line.

Preparing C programs in HP-UX:

This topic contains information to consider when preparing C programs in HP-UX; with examples for the IA64 (IPF) platform.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

Work in your normal environment. Precompiled C programs are supplied in the *MQ_INSTALLATION_PATH/samp/bin* directory.

For further information about programming 64 bit applications see *Coding standards on 64-bit platforms*.

To use TLS, IBM MQ MQI clients on HP-UX must be built using POSIX threads.

Some examples to consider are:

- “IA64 (IPF) platform”
- “Linking libraries” on page 395

IA64 (IPF) platform

Build examples of *amqsput0*, *cliexit*, and *srvexit* on IA64(IPF) platform.

The following example builds the sample program *amqsput0* as a client application in a non-threaded 32 bit environment:

```

c89 -w1,+b,:+e -D HPUX_SOURCE -o amqsputc_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -lmqic

```

The following example builds the sample program *amqsput0* as a client application in a threaded 32 bit environment:

```

c89 -mt -w1,+b,:+e -D HPUX_SOURCE -o amqsputc_32_r amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -lmqic_r -lpthread

```

The following example builds the sample program *amqsput0* as a client application in a non-threaded 64 bit environment:

```

c89 +DD64 +e -D HPUX_SOURCE -o amqsputc_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqic

```

The following example builds the sample program amqsput0 as a client application in a threaded 64 bit environment:

```
c89 -mt +DD64 +e -D_HPUX_SOURCE -o amqsputc_64_r amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqic_r -lpthread
```

The following example builds the sample program amqsput0 as a server application in a non-threaded 32 bit environment:

```
c89 -Wl,+b,: +e -D_HPUX_SOURCE -o amqsput_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -lmqm
```

The following example builds the sample program amqsput0 as a server application in a threaded 32 bit environment:

```
c89 -mt -Wl,+b,: +e -D_HPUX_SOURCE -o amqsput_32_r amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -lmqm_r -lpthread
```

The following example builds the sample program amqsput0 as a server application in a non-threaded 64 bit environment:

```
c89 +DD64 +e -D_HPUX_SOURCE -o amqsput_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqm
```

The following example builds the sample program amqsput0 as a server application in a threaded 64 bit environment:

```
c89 -mt +DD64 +e -D_HPUX_SOURCE -o amqsput_64_r amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqm_r -lpthread
```

The following example builds a client exit cliexit in a non-threaded 32 bit environment:

```
c89 +e +z -c -D_HPUX_SOURCE -o cliexit.o cliexit.c -I MQ_INSTALLATION_PATH/inc
ld +b: -b cliexit.o +ee MQStart -o /var/mqm/exits/cliexit_32 -L MQ_INSTALLATION_PATH/lib \
-L/usr/lib/hpux32 -lmqic
```

The following example builds a client exit cliexit in a threaded 32 bit environment:

```
c89 -mt +e +z -c -D_HPUX_SOURCE -o cliexit.o cliexit.c -I MQ_INSTALLATION_PATH/inc
ld +b: -b cliexit.o +ee MQStart -o /var/mqm/exits/cliexit_32_r -L MQ_INSTALLATION_PATH/lib \
-L/usr/lib/hpux32 -lmqic_r -lpthread
```

The following example builds a client exit cliexit in a non-threaded 64 bit environment:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o cliexit.o cliexit.c -I MQ_INSTALLATION_PATH/inc
ld -b cliexit.o +ee MQStart -o /var/mqm/exits64/cliexit_64 \
-L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqic
```

The following example builds a client exit cliexit in a threaded 64 bit environment:

```
c89 -mt +DD64 +e +z -c -D_HPUX_SOURCE -o cliexit.o cliexit.c -I MQ_INSTALLATION_PATH/inc
ld -b cliexit.o +ee MQStart -o /var/mqm/exits/cliexit_64_r \
-L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqic_r -lpthread
```

The following example builds a server exit srvexit in a non-threaded 32 bit environment:

```
c89 +e +z -c -D_HPUX_SOURCE -o srvexit.o srvexit.c -I MQ_INSTALLATION_PATH/inc
ld +b: -b srvexit.o +ee MQStart -o /var/mqm/exits/srvexit_32 -L MQ_INSTALLATION_PATH/lib \
-L/usr/lib/hpux32 -lmqm
```

The following example builds a server exit srvexit in a threaded 32 bit environment:

```
c89 -mt +e +z -c -D_HPUX_SOURCE -o srvexit.o srvexit.c -I MQ_INSTALLATION_PATH/inc
ld +b: -b srvexit.o +ee MQStart -o /var/mqm/exits/srvexit_32_r -L MQ_INSTALLATION_PATH/lib \
-L/usr/lib/hpux32 -lmqm_r -lpthread
```

The following example builds a server exit srvexit in a non-threaded 64 bit environment:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o srvexit.o srvexit.c -I MQ_INSTALLATION_PATH MQ_INSTALLATION_PATH/inc
ld -b srvexit.o +ee MQStart -o /var/mqm/exits64/srvexit_64 \
-L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqm
```

The following example builds a server exit `srvexit` in a threaded 64 bit environment:

```
c89 -mt +DD64 +e +z -c -D_HPUX_SOURCE -o srvexit.o srvexit.c -I MQ_INSTALLATION_PATH/inc
ld -b srvexit.o +ee MQStart -o /var/mqm/exits/srvexit_64_r \
-L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqm_r -lpthread
```

Linking libraries

You need to link your programs with one of the libraries provided by IBM MQ.

The following table shows which library to use in different environments:

Hardware platform	Threaded or non-threaded environment	Program/exit type	Library file
IA64 (IPF)	Threaded	Server & Client for C	libmqm_r.so
IA64 (IPF)	Threaded	Client for C	libmqic_r.so
IA64 (IPF)	Non-threaded	Server & Client for C	libmqm.so
IA64 (IPF)	Non-threaded	Client for C	libmqic.so

Note:

1. You cannot link to more than one library. That is, you cannot link to both a threaded and a non-threaded library at the same time.
2. If you are writing an installable service (see the Administering for further information), you need to link to the `libmqmzf.sl` library.
3. If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries Encina, or BEA Tuxedo, you need to link to the `libmqmxa.sl` (or `libmqmxa64.sl` if your transaction manager treats the 'long' type as 64 bit) and `libmqz.sl` libraries in a non-threaded application and to the `libmqmxa_r.sl` (or `libmqmxa64_r.sl`) and `libmqz_r.sl` libraries in a threaded application.
4. You must link IBM MQ libraries before any other product libraries.

Preparing COBOL programs in HP-UX:

Learn about preparing COBOL programs in HP-UX, using Micro Focus Server Express with IBM MQ on the IA64 (IPF) platform, and running programs in the IBM MQ MQI client environment.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Usage notes:

1. 32 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

and symbolic links are created in:

```
MQ_INSTALLATION_PATH/inc
```

2. 64 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

3. In the following examples set COBCPY to:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

for 32 bit applications, and:
`MQ_INSTALLATION_PATH/inc/cobcpy64`

for 64 bit applications.

Compile the programs using the Micro Focus compiler. The copy files that declare the structures are in `MQ_INSTALLATION_PATH/inc`:

```
$ export LIB=MQ_INSTALLATION_PATH/lib:$LIB  
$ export COBCPY="COBCPY_VALUE"
```

Compiling 32 bit programs:

```
$ cob32 -xv amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqmc Server for COBOL  
$ cob32 -xv amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqic Client for COBOL  
$ cob32 -xtv amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqmc_r Threaded Server for COBOL  
$ cob32 -xtv amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqic_r Threaded Client for COBOL
```

Compiling 64 bit programs:

```
$ cob64 -xv amqsput.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmc Server for COBOL  
$ cob64 -xv amqsput.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqic Client for COBOL  
$ cob64 -xtv amqsput.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmc_r Threaded Server for COBOL  
$ cob64 -xtv amqsput.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqic_r Threaded Client for COBOL
```

where `amqsput` is a sample program

Ensure that you have specified adequate runtime stack sizes; 16 KB is the recommended minimum.

You need to link your programs with the appropriate library provided by IBM MQ. The following table shows which library to use in different environments

Hardware platform	Program/exit type	Library file
IA64 (IPF)	Server for COBOL	libmqmc.so
IA64 (IPF)	Client for COBOL	libmqic.so
IA64 (IPF)	Threaded applications	libmqmc_r.so

Using Micro Focus Server Express with IBM MQ on the IA64 (IPF) platform

See “Address Space models supported by IBM MQ for HP-UX on IA64 (IPF)” on page 398 for details on using Micro Focus Server Express in conjunction with IBM MQ on the HP/IPF platform.

Programs to run in the IBM MQ MQI client environment

If you are using LU 6.2 to connect your MQI client to a server, link your application to `libsna.a`, part of the SNAplusAPI product. Use the `-lv3` and `-lstr` options on your compiling and linking command.

- The `-lv3` option gives your program access to the AT&T signaling library (the SNAplusAPI uses AT&T signals)
- The `-lstr` option links your program to the streams component

Preparing CICS programs in HP-UX:

Learn to build CICS transaction programs in HP-UX.

To build the sample CICS transaction, amqscic0.ccs, run the following command:

```
$ export USERLIB="-lmqm_r"  
$ cicstcl -l C amqscic0.ccs
```

An XA switch module is provided to enable you to link CICS with IBM MQ:

Table 54. Essential code for CICS applications (HP-UX)

Description	C (source)	C (exec)
XA initialization routine	amqzscix.c	amqzsc

You can find more information about supporting CICS transactions in the Administering.

TXSeries CICS support:

IBM MQ on HP-UX supports TXSeries CICS using the XA interface. Ensure that CICS applications are linked to the threaded version of the MQ libraries.

Write IBM MQ programs that are loaded into the same CICS region in either C or COBOL. You cannot make a combination of C and COBOL MQI calls into the same CICS region. Most MQI calls in the second language used fail with a reason code of MQRC_HOBBJ_ERROR.

CICS C sample transaction

Sample C source for a CICS IBM MQ transaction is provided by AMQSCIC0.CCS. The transaction reads messages from the transmission queue SYSTEM.SAMPLE.CICS.WORKQUEUE on the default queue manager and places them onto the local queue with the queue name that is contained in the transmission header of the message. Any failures are sent to the queue SYSTEM.SAMPLE.CICS.DLQ. Use the sample MQSC script AMQSCIC0.TST to create these queues and sample input queues.

Preparing CICS COBOL programs using Micro Focus COBOL

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

To use Micro Focus COBOL, follow these steps:

1. Add the IBM MQ COBOL runtime library module to the runtime library using the following command:

```
cicsmkcobol -L/usr/lib/dce -L MQ_INSTALLATION_PATH/lib \  
            MQ_INSTALLATION_PATH/lib/libmqmcbt.o -lmqe_r
```

Note: With `cicsmkcobol`, IBM MQ does not allow you to make MQI calls in the C programming language from your COBOL application.

If your existing applications have any such calls, you are recommended to move these functions from the COBOL applications to your own library, for example, `myMQ.so`. After moving these functions do not include the IBM MQ library `libmqmcbt.o` when building the COBOL application for CICS.

Additionally, if your COBOL application does not make any COBOL MQI call, do not link `libmqmz_r` with `cicsmkcobol`.

This creates the Micro Focus COBOL language method file and enables the CICS runtime COBOL library to call IBM MQ on UNIX and Linux systems.

Note: Run `cicsmkcobol` only when you install one of the following products:

- New version or release of Micro Focus COBOL
- New version or release of CICS for HP-UX
- New version or release of any supported database product (for COBOL transactions only)
- New version or release of IBM MQ

2. Export the following environment variable:

```
COBCPY= MQ_INSTALLATION_PATH/inc export COBCPY
```

3. Translate, compile, and link the program by typing:

```
cicstcl -l COBOL -e yourprog.ccp
```

Address Space models supported by IBM MQ for HP-UX on IA64 (IPF):

The HP-UX provides several address space models that can be exploited by IBM MQ applications.

HP-UX supports two Address Space models:

- MGAS - Mostly Global Address space (this is the default and is used by IBM MQ)
- MPAS - Mostly Private Address space

Applications which connect to IBM MQ can use either the MGAS or MPAS address space models. Applications built using the MPAS model that connect to IBM MQ using shared memory might incur a minor performance cost due to the inefficiency in mapping the shared memory pages used by IBM MQ into the virtual address space of the MPAS program.

COBOL applications built using Micro Focus Server Express use the MPAS model by default.

You can use the **chatr** program to check and change the addressing model used by a program.

If you encounter problems connecting to IBM MQ from 32-bit MPAS programs, consider using the MGAS addressing model, or building your application as a 64-bit MPAS application rather than a 32-bit MPAS application.

More details on the MGAS and MPAS address space models can be found in the HP-UX documentation.

Building your procedural application on Linux



This information describes the additional tasks, and the changes to the standard tasks, that you must perform when building IBM MQ for Linux applications to run.

C and C++ are supported. For information about preparing your C++ programs, see Using C++.

Preparing C programs in Linux:

Precompiled C programs are supplied in the `MQ_INSTALLATION_PATH/samp/bin` directory. To build a sample from source code, use the `gcc` compiler.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Work in your normal environment. For further information about programming 64 bit applications, see Coding standards on 64-bit platforms.

Linking libraries

The following tables lists the libraries that are needed when preparing C programs on Linux.

- You need to link your programs with the appropriate library provided by IBM MQ.

In a non-threaded environment, link to only one of the following libraries:

Library file	Program/exit type
<code>libmqm.so</code>	Server for C
<code>libmqic.so</code> & <code>libmqm.so</code>	Client for C

In a threaded environment, link to only one of the following libraries:

Library file	Program/exit type
<code>libmqm_r.so</code>	Server for C
<code>libmqic_r.so</code> & <code>libmqm_r.so</code>	Client for C

Note:

1. You cannot link to more than one library. That is, you cannot link to both a threaded and a non-threaded library at the same time.
2. If you are writing an installable service (see the Administering for further information), you need to link to the `libmqmf.so` library.
3. If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries Encina, or BEA Tuxedo, you need to link to the `libmqmxa.so` (or `libmqmxa64.so` if your transaction manager treats the 'long' type as 64 bit) and `libmqz.so` libraries in a non-threaded application and to the `libmqmxa_r.so` (or `libmqmxa64_r.so`) and `libmqz_r.so` libraries in a threaded application.
4. You must link IBM MQ libraries before any other product libraries.

Building 31-bit applications:

This topic contains examples of the commands used to build 31-bit programs in various environments.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

C client application, 31-bit, non-threaded

```
gcc -m31 -o famqsputc_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqic
```

C client application, 31-bit, threaded

```
gcc -m31 -o amqsputc_32_r amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqic_r -lpthread
```

C server application, 31-bit, non-threaded

```
gcc -m31 -o amqsput_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm
```

C server application, 31-bit, threaded

```
gcc -m31 -o amqsput_32_r amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm_r -lpthread
```

C++ client application, 31-bit, non-threaded

```
g++ -m31 -fsigned-char -o imqsputc_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqc23gl  
-limqb23gl -lmqic
```

C++ client application, 31-bit, threaded

```
g++ -m31 -fsigned-char -o imqsputc_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqc23gl_r  
-limqb23gl_r -lmqic_r -lpthread
```

C++ server application, 31-bit, non-threaded

```
g++ -m31 -fsigned-char -o imqsput_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqs23gl  
-limqb23gl -lmqm
```

C++ server application, 31-bit, threaded

```
g++ -m31 -fsigned-char -o imqsput_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqs23gl_r  
-limqb23gl_r -lmqm_r -lpthread
```

C client exit, 31-bit, non-threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/cliexit_32 cliexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -lmqic
```

C client exit, 31-bit, threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/cliexit_32_r cliexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -lmqic_r -lpthread
```

C server exit, 31-bit, non-threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/srvexit_32 srvexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -lmqm
```

C server exit, 31-bit, threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/srvexit_32_r srvexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -lmqm_r -lpthread
```


Building 32-bit applications:

This topic contains examples of the commands used to build 32-bit programs in various environments.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

C client application, 32-bit, non-threaded

```
gcc -m32 -o amqsputc_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqic
```

C client application, 32-bit, threaded

```
gcc -m32 -o amqsputc_32_r amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqic_r -lpthread
```

C server application, 32-bit, non-threaded

```
gcc -m32 -o amqsput_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm
```

C server application, 32-bit, threaded

```
gcc -m32 -o amqsput_32_r amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm_r -lpthread
```

C++ client application, 32-bit, non-threaded

```
g++ -m32 -fsigned-char -o imqsputc_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl -limqb23gl -lmqic
```

C++ client application, 32-bit, threaded

```
g++ -m32 -fsigned-char -o imqsputc_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl_r -limqb23gl_r -lmqic_r -lpthread
```

C++ server application, 32-bit, non-threaded

```
g++ -m32 -fsigned-char -o imqsput_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqs23gl -limqb23gl -lmqm
```

C++ server application, 32-bit, threaded

```
g++ -m32 -fsigned-char -o imqsput_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqs23gl_r -limqb23gl_r -lmqm_r -lpthread
```

C client exit, 32-bit, non-threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/cliexit_32 cliexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -lmqic
```

C client exit, 32-bit, threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/cliexit_32_r cliexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -lmqic_r -lpthread
```

C server exit, 32-bit, non-threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/srvexit_32 srvexit.c -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -lmqm
```

C server exit, 32-bit, threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/srvexit_32_r srvexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -lmqm_r -lpthread
```

Building 64-bit applications:

This topic contains examples of the commands used to build 64-bit programs in various environments.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

C client application, 64-bit, non-threaded

```
gcc -m64 -o amqsputc_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqic
```

C client application, 64-bit, threaded

```
gcc -m64 -o amqsputc_64_r amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqic_r
-lpthread
```

C server application, 64-bit, non-threaded

```
gcc -m64 -o amqsput_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqm
```

C server application, 64-bit, threaded

```
gcc -m64 -o amqsput_64_r amqsput0.c -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqm_r
-lpthread
```

C++ client application, 64-bit, non-threaded

```
g++ -m64 -fsigned-char -o imqsputc_64 imqsput.cpp
-I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-lmqc23gl -limqb23gl -lmqic
```

C++ client application, 64-bit, threaded

```
g++ -m64 -fsigned-char -o imqsputc_64_r imqsput.cpp
-I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-lmqc23gl_r -limqb23gl_r -lmqic_r -lpthread
```

C++ server application, 64-bit, non-threaded

```
g++ -m64 -fsigned-char -o imqsput_64 imqsput.cpp
-I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqs23gl -limqb23gl -lmqm
```

C++ server application, 64-bit, threaded

```
g++ -m64 -fsigned-char -o imqsput_64_r imqsput.cpp
-I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqs23gl_r -limqb23gl_r -lmqm_r -lpthread
```

C client exit, 64-bit, non-threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/cliexit_64 cliexit.c
-I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqic
```

C client exit, 64-bit, threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/cliexit_64_r cliexit.c
-I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqic_r -lpthread
```

C server exit, 64-bit, non-threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/srvexit_64 srvexit.c
-I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqm
```

C server exit, 64-bit, threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/srvexit_64_r srvexit.c
-I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqm_r -lpthread
```

Preparing COBOL programs in Linux:

Learn about preparing COBOL programs in Linux and preparing COBOL programs using Micro Focus COBOL.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

1. 32 bit COBOL copy books are installed in the following directory:

MQ_INSTALLATION_PATH/inc/cobcpy32

and symbolic links are created in:

MQ_INSTALLATION_PATH/inc

2. On 64 bit platforms, 64 bit COBOL copy books are installed in the following directory:

MQ_INSTALLATION_PATH/inc/cobcpy64

3. In the following examples set COBCPY to:

MQ_INSTALLATION_PATH/inc/cobcpy32

for 32 bit applications, and:

MQ_INSTALLATION_PATH/inc/cobcpy64

for 64 bit applications.

You need to link your program with one of the following:

Library file	Program/exit type
libmqmcb.so	Server for COBOL
libmqicb.so	Client for COBOL
libmqmcb_r.so	Server for COBOL (threaded application)
libmqicb_r.so	Client for COBOL (threaded application)

Preparing COBOL programs using Micro Focus COBOL

Set environment variables before compiling your program as follows:

```
export COBCPY=COBCPY_VALUE
export LIB=MQ_INSTALLATION_PATH lib:$LIB
```

To compile a 32 bit COBOL program, where supported, using Micro Focus COBOL, enter:

```
$ cob32 -xvP amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqcb Server for COBOL
$ cob32 -xvP amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb Client for COBOL
$ cob32 -xtvP amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqcb_r Threaded Server for COBOL
$ cob32 -xtvP amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb_r Threaded Client for COBOL
```

To compile a 64 bit COBOL program using Micro Focus COBOL, enter:

```
$ cob64 -xvP amqsput.cb1 -L MQ_INSTALLATION_PATH/lib64 -lmqmb Server for COBOL
$ cob64 -xvP amqsput.cb1 -L MQ_INSTALLATION_PATH/lib64 -lmqicb Client for COBOL
$ cob64 -xtvP amqsput.cb1 -L MQ_INSTALLATION_PATH/lib64 -lmqmb_r Threaded Server for COBOL
$ cob64 -xtvP amqsput.cb1 -L MQ_INSTALLATION_PATH/lib64 -lmqicb_r Threaded Client for COBOL
```

where amqsput is a sample program

See the Micro Focus COBOL documentation for a description of the environment variables that you need.

Building your procedural application on IBM i



The IBM i publications describe how to build executable applications from the programs that you write, to run with IBM i on iSeries or System i® systems.

This topic describes the additional tasks, and the changes to the standard tasks, that you must perform when building IBM MQ for IBM i procedural applications to run on IBM i systems. COBOL, C, C++, Java and RPG programming languages are supported. For information about preparing your C++ programs, see Using C++. For information about preparing your Java programs, see Using IBM MQ classes for Java.

The tasks that you must perform to create an executable IBM MQ for IBM i application depend on the programming language that the source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the IBM MQ for IBM i data definition files for the language that you are using. Make yourself familiar with the contents of these files. See “IBM MQ data definition files” on page 71 for a full description.

Preparing C programs in IBM i:

IBM MQ for IBM i supports messages up to 100 MB in size. Application programs written in ILE C, supporting IBM MQ messages greater than 16 MB, need to use the *Teraspace* compiler option to allocate sufficient memory for these messages.

For more information about the C compiler options, see the *WebSphere Development Studio ILE C/C++ Programmer's Guide*.

To compile a C module, you can use the IBM i command, CRTCMOD. Make sure that the library containing the include files (QMQM) is in the library list when you compile.

You must then bind the output of the compiler with the service program using the CRTPGM command.

An example of the command for a nonthreaded environment is:

Table 55. Example of CRTPGM in the nonthreaded environment

Command	Program/exit type
CRTPGM PGM(<i>pgmname</i>) MODULE(<i>pgmname</i>) BNDSRVPGM(QMQM/LIBMQM)	Server or client for C

where *pgmname* is the name of your program.

An example of the command for a threaded environment is:

Table 56. Example of CRTPGM in the threaded environment

Command	Program/exit type
CRTPGM PGM(<i>pgmname</i>) MODULE(<i>pgmname</i>) BNDSRVPGM(QMQM/LIBMQM_R)	Server or client for C

where *pgmname* is the name of your program.

The following tables list the libraries that are needed when preparing C programs on IBM i in a non-threaded environment and threaded environment.

Table 57. Non-threaded environment

Library file	Program/exit type
LIBMQM	Server for C
LIBMQIC & LIBMQM	Client for C

Table 58. Threaded environment

Library file	Program/exit type
LIBMQM_R	Server for C
LIBMQIC_R & LIBMQM_R	Client for C

Preparing COBOL programs in IBM i:

Learn about preparing COBOL programs in IBM i and the two methods of accessing the MQI from within the COBOL program.

IBM MQ for IBM i provides two methods for accessing the MQI from within COBOL programs:

1. A dynamic call interface to programs having the names of the MQI functions, such as MQCONN and MQOPEN. This interface is intended primarily for use with the OPM (Original Program Mode) COBOL compiler, but can also be used with the ILE (Integrated Language Environment) COBOL compiler. Some functions in IBM MQ for IBM i, such as MQCMIT and MQBACK, are not supported through this interface, which is provided for compatibility with previous releases.
2. A bound procedural call interface provided by service programs. This provides access to all the MQI functions in IBM MQ for IBM i, support for threaded applications, and potentially better performance than the dynamic call interface. This interface can be used only with the ILE COBOL compiler.

In both cases the standard COBOL CALL syntax is used to access the MQI functions.

The COBOL copy files containing the named constants and structure definitions for use with the MQI are contained in the source physical file QMQM/QCBLLESRC.

The COBOL copy files use the single quotation mark character (') as the string delimiter. The IBM i COBOL compilers assume that the delimiter is the quotation mark ("). To prevent the compilers generating warning messages, specify OPTION(*APOST) on the commands CRTCBPLPGM, CRTBNDCBL, or CRTCBLMOD.

To make the compiler accept the single quotation mark character (') as the string delimiter in the COBOL copy files, use the compiler option \APOST.

Using the dynamic call interface

- The QMQM library must be in your library list when you compile and when you run COBOL programs using the MQI dynamic call interface.

- Use the CRTCBPLPGM command to invoke the OPM COBOL compiler.
- Use either the CRTBNDCBL command or the two separate commands CRTCBMOD and CRTPGM to invoke the ILE COBOL compiler.

Using the bound procedure call interface

- First create a module using the CRTCBMOD compiler specifying the parameter:
LINKLIT(*PRC)
- Then use the CRTPGM command to create the program object specifying the parameter:

for non-threaded applications

```
BNDSRVPGM(QMQM/AMQ0STUB) Server for COBOL for non-threaded applications
BNDSRVPGM(QMQM/AMQCSTUB) Client for COBOL for non-threaded applications
```

for threaded applications

```
BNDSRVPGM(QMQM/AMQ0STUB_R) Server for COBOL for threaded applications
BNDSRVPGM(QMQM/AMQCSTUB_R) Client for COBOL for threaded applications
```

Note: Except for programs created using the V4R4 ILE COBOL compiler and containing the THREAD(SERIALIZE) option in the PROCESS statement, COBOL programs must not use the threaded IBM MQ libraries. Even if a COBOL program has been made thread safe in this manner, be careful when you design the application, because THREAD(SERIALIZE) forces serialization of COBOL procedures at the module level and might affect overall performance.

See the *WebSphere Development Studio: ILE COBOL Programmer's Guide* and the *WebSphere Development Studio: ILE COBOL Reference* for further information.

For more information about compiling a CICS application, see the *CICS for IBM i Application Programming Guide*, SC41-5454.

Preparing CICS programs in IBM i:

Learn about the steps required when preparing CICS programs in IBM i.

To create a program that includes EXEC CICS statements and MQI calls, perform these steps:

1. If necessary, prepare maps using the CRTICSMAP command.
2. Translate the EXEC CICS commands into native language statements. Use the CRTICSC command for a C program. Use the CRTICSCBL command for a COBOL program.
Include CICSOPT(*NOGEN) in the CRTICSC or CRTICSCBL command. This halts processing to enable you to include the appropriate CICS and IBM MQ service programs. This command puts the code, by default, into QTEMP/QACYCICS.
3. Compile the source code using the CRTCMOD command (for a C program) or the CRTCBMOD command (for a COBOL program).
4. Use CRTPGM to link the compiled code with the appropriate CICS and IBM MQ service programs. This creates the executable program.

An example of such code follows (it compiles the shipped CICS sample program):

```
CRTICSC OBJ(QTEMP/AMQSCIC0) SRCFILE(/MQSAMP/QCSRC) +
  SRCMBR(AMQSCIC0) OUTPUT(*PRINT) +
  CICSOPT(*SOURCE *NOGEN)
CRTCMOD MODULE(MQTEST/AMQSCIC0) +
  SRCFILE(QTEMP/QACYCICS) OUTPUT(*PRINT)
CRTPGM PGM(MQTEST/AMQSCIC0) MODULE(MQTEST/AMQSCIC0) +
  BNDSRVPGM(QMQM/LIBMQC QCICS/AEGEIPGM)
```

Preparing RPG programs in IBM i:

If you are using IBM MQ for IBM i, you can write your applications in RPG.

For more information see “Coding IBM MQ programs in RPG (IBM i only)” on page 455, and refer to the IBM i Application Programming Reference (ILE/RPG).

SQL programming considerations:

Learn about the steps required when building an application on IBM i using SQL.

If your program contains EXEC SQL statements and MQI calls, perform these steps:

1. Translate the EXEC SQL commands into native language statements. Use the CRTSQLCI command for a C program. Use the CRTSQLCBLI command for a COBOL program.
Include OPTION(*NOGEN) in the CRTSQLCI or CRTSQLCBLI command. This halts processing to enable you to include the appropriate IBM MQ service programs. This command puts the code, by default, into QTEMP/QSQLTEMP.
2. Compile the source code using the CRTCMOD command (for a C program) or the CRTCBMOD command (for a COBOL program).
3. Use CRTPGM to link the compiled code with the appropriate IBM MQ service programs. This creates the executable program.

An example of such code follows (it compiles a program, SQLTEST, in library, SQLUSER):

```
CRTSQLCI OBJ(MQTEST/SQLTEST) SRCFILE(SQLUSER/QCSRC) +  
          SRCMBR(SQLTEST) OUTPUT(*PRINT) OPTION(*NOGEN)  
CRTCMOD  MODULE(MQTEST/SQLTEST) +  
          SRCFILE(QTEMP/QSQLTEMP) OUTPUT(*PRINT)  
CRTPGM  PGM(MQTEST/SQLTEST) +  
          BNDSRVPGM(QMQM/LIBMQIC)
```

Building your procedural application on Solaris



This information describes the additional tasks, and the changes to the standard tasks, that you must perform when building IBM MQ for Solaris applications to run under Solaris.

COBOL, C, and C++ programming languages are supported. For information about preparing your C++ programs, see Using C++.

In addition to coding the MQI calls in your source code, you must add the appropriate include files. Make yourself familiar with the contents of these files. See “IBM MQ data definition files” on page 71 for a full description.

Throughout this topic, the backslash (\) character is used to split long commands over more than one line. Do not enter this character, enter each command as a single line.

Preparing C programs in Solaris:

Precompiled C programs are supplied in the `MQ_INSTALLATION_PATH/samp/bin` directory.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

For further information about programming 64 bit applications, see Coding standards on 64-bit platforms.

If you want to use the programs on a machine that has only the IBM MQ MQI client for Solaris installed, compile the programs to link them with the client library (`-lmqic`).

If you use the unsupported compiler `/usr/ucb/cc`, your application might compile and link successfully. However, when you run the application, it fails when it attempts to connect to the queue manager.

Note: 32 bit Solaris x86 SSL and TLS clients configured for FIPS 140-2 compliant operation fail when running on Intel systems. This failure occurs because the FIPS 140-2 compliant GSKit-Crypto Solaris x86 32 bit library file does not load on the Intel chipset. On affected systems, error AMQ9655 is reported in the client error log. To resolve this issue, disable FIPS 140-2 compliance or recompile the client application 64 bit, because 64 bit code is not affected.

Linking libraries

You must link with the IBM MQ libraries that are appropriate for your application type:

Library files	Program/exit type
libmqm.so	Server for C
libmqic.so & libmqm.so	Client for C

Note:

1. If you are writing an installable service (for further information, see the Administering), link to the `libmqmf.so` library.
2. If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries Encina, or BEA Tuxedo, you must link to the `libmqmxa.so` (or `libmqmxa64.so` if your transaction manager treats the 'long' type as 64 bit) and `libmqz.so` libraries.
3. You must link IBM MQ libraries before any other product libraries.

Building applications on x86-64:

This topic contains examples of the commands used to build programs in various environments on the x86-64 platform.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

C client application, 32-bit

```
cc -xarch=386 -mt -o amqsputc_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -lmqic -lsocket -lnsl -ldl
```

C client application, 64-bit

```
cc -xarch=amd64 -mt -o amqsputc_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqic -lsocket -lnsl -ldl
```

C server application, 32-bit

```
cc -xarch=386 -mt -o amqsput_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -lmqm -lsocket -lnsl -ldl
```


C server application, 64-bit

```
cc -xarch=amd64 -mt -o amqsput_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqm -lsocket -lnsl -ldl
```

C++ client application, 32-bit

```
CC -xarch=386 -mt -o imqsputc_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqc23as -limqb23as -lmqic -lsocket -lnsl -ldl
```

C++ client application, 64-bit

```
CC -xarch=amd64 -mt -o imqsputc_64 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqc23as -limqb23as  
-lmqic -lsocket -lnsl -ldl
```

C++ server application, 32-bit

```
CC -xarch=386 -mt -o imqsput_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqs23as -limqb23as -lmqm  
-lsocket -lnsl -ldl
```

C++ server application, 64-bit

```
CC -xarch=amd64 -mt -o imqsput_64 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqs23as -limqb23as -lmqm  
-lsocket -lnsl -ldl
```

C client exit, 32-bit

```
cc -xarch=386 -mt -G -KPIC -o /var/mqm/exits/cliexit_32 cliexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32  
-lmqic -lsocket -lnsl -ldl
```

C client exit, 64-bit

```
cc -xarch=amd64 -mt -G -KPIC -o /var/mqm/exits64/cliexit_64 cliexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64  
-lmqic -lsocket -lnsl -ldl
```

C server exit, 32-bit

```
cc -xarch=386 -mt -G -KPIC -o /var/mqm/exits/srvexit_32 srvexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32  
-lmqm -lsocket -lnsl -ldl
```

C server exit, 64-bit

```
cc -xarch=amd64 -mt -G -KPIC -o /var/mqm/exits64/srvexit_64 srvexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64  
-lmqm -lsocket -lnsl -ldl
```

Building applications on SPARC:

This topic contains examples of the commands used to build programs in various environments on the SPARC platform.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

C client application, 32-bit

```
cc -xarch=v8plus -mt -o amqsputc_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -lmqic -lsocket -lnsl -ldl
```

C client application, 64-bit

```
cc -xarch=v9 -mt -o amqsputc_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqic  
-lsocket -lnsl -ldl
```

C server application, 32-bit

```
cc -xarch=v8plus -mt -o amqsput_32 amqsput0.c -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -lmqm -lsocket -lnsl -ldl
```

C server application, 64-bit

```
cc -xarch=v9 -mt -o amqsput_64 amqsput0.c -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqm  
-lsocket -lnsl -ldl
```

C++ client application, 32-bit

```
CC -xarch=v8plus -mt -o imqsputc_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqc23as -limqb23as -lmqic  
-lsocket -lnsl -ldl
```

C++ client application, 64-bit

```
CC -xarch=v9 -mt -o imqsputc_64 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqc23as -limqb23as  
-lmqic -lsocket -lnsl -ldl
```

C++ server application, 32-bit

```
CC -xarch=v8plus -mt -o imqsput_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib  
-R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqs23as -limqb23as -lmqm  
-lsocket -lnsl -ldl
```

C++ server application, 64-bit

```
CC -xarch=v9 -mt -o imqsput_64 imqsput.cpp -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqs23as -limqb23as -lmqm  
-lsocket -lnsl -ldl
```

C client exit, 32-bit

```
cc -xarch=v8plus -mt -G -KPIC -o /var/mqm/exits/cliexit_32 cliexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32  
-lmqic -lsocket -lnsl -ldl
```

C client exit, 64-bit

```
cc -xarch=v9 -mt -G -KPIC -o /var/mqm/exits64/cliexit_64 cliexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64  
-lmqic -lsocket -lnsl -ldl
```

C server exit, 32-bit

```
cc -xarch=v8plus -mt -G -KPIC -o /var/mqm/exits/srvexit_32 srvexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32  
-lmqm -lsocket -lnsl -ldl
```

C server exit, 64-bit

```
cc -xarch=v9 -mt -G -KPIC -o /var/mqm/exits64/srvexit_64 srvexit.c  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64  
-lmqm -lsocket -lnsl -ldl
```

Preparing COBOL programs in Solaris:

Learn about preparing COBOL programs in Solaris.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

1. 32 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

and symbolic links are created in:

```
MQ_INSTALLATION_PATH/inc
```

2. 64 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

3. In the following examples set COBCPY to:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

for 32 bit applications, and:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

for 64 bit applications.

Compile the programs using Micro Focus compiler. The copy files that declare the structures are in `MQ_INSTALLATION_PATH/inc`:

```
$ export LIB= MQ_INSTALLATION_PATH/lib:$LIB
$ export COBCPY="COBCPY_VALUE"
```

Compiling 32 bit programs:

- `$ cob32 -xv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqmc`
Server for COBOL
- `$ cob32 -xv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb`
Client for COBOL
- `$ cob32 -xtv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqmc_r`
Threaded Server for COBOL
- `$ cob32 -xtv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb_r`
Threaded Client for COBOL

Compiling 64-bit programs:

- `$ cob64 -xv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmc`
Server for COBOL
- `$ cob64 -xv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb`
Client for COBOL
- `$ cob64 -xtv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmc_r`
Threaded Server for COBOL
- `$ cob64 -xtv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb_r`
Threaded Client for COBOL

where `amqs0put0.cbl` is a sample program.

You must link your program with one of the following:

- `libmqmc.so`
Server for COBOL
- `libmqicb.so`
Client for COBOL

Preparing CICS programs in Solaris:

Learn about preparing CICS programs in Solaris.

An XA switch module is provided to enable you to link CICS with IBM MQ:

Table 59. Essential code for CICS applications (Solaris)

Description	C (source)	C (exec)
XA initialization routine	amqzscix.c	amqzsc - TXSeries for Solaris

Always link your transactions with the thread safe IBM MQ library `libmqm.so`.

You can find more information about supporting CICS transactions in the Administering.

TXSeries CICS support:

IBM MQ for Solaris supports TXSeries CICS using the XA interface.

Write IBM MQ programs that are loaded into the same CICS region in either C or COBOL. You cannot make a combination of C and COBOL MQI calls into the same CICS region. Most MQI calls in the second language used fail with a reason code of MQRC_HOBBJ_ERROR.

Preparing CICS COBOL programs using Micro Focus COBOL

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

To use Micro Focus COBOL, follow these steps:

1. Add the IBM MQ COBOL runtime library module to the runtime library using the following command:

```
cicsmkcobol -L/usr/lib/dce -L MQ_INSTALLATION_PATH/lib \  
            MQ_INSTALLATION_PATH/lib/libmqmcbt.o -lmqe
```

Note: With `cicsmkcobol`, IBM MQ does not allow you to make MQI calls in the C programming language from your COBOL application.

If your existing applications have any such calls, move these functions from the COBOL applications to your own library, for example, `myMQ.so`. After moving these functions do not include the IBM MQ library `libmqmcbt.o` when building the COBOL application for CICS.

Additionally, if your COBOL application does not make any COBOL MQI call, do not link `libmqmz_r` with `cicsmkcobol`.

This creates the Micro Focus COBOL language method file and enables the CICS runtime COBOL library to call IBM MQ on UNIX and Linux systems.

Note: Run `cicsmkcobol` only when you install one of the following products:

- New version or release of Micro Focus COBOL
 - New version or release of TXSeries for Solaris
 - New version or release of any supported database product (for COBOL transactions only)
 - New version or release of IBM MQ
2. Export the following environment variable:
`COBCPY= MQ_INSTALLATION_PATH/inc export COBCPY`
 3. Translate, compile, and link the program by typing:

```
cicstcl -l COBOL -e yourprog.ccp
```

Preparing CICS C programs

Build CICS C programs using the standard CICS facilities:

1. Export *one* of the following environment variables:
 - `LDFLAGS = "-L MQ_INSTALLATION_PATH ∆lib -lmqm_r" export LDFLAGS`
 - `USERLIB = "-L MQ_INSTALLATION_PATH ∆lib -lmqm_r" export USERLIB`
2. Translate, compile, and link the program by typing:

```
cicstcl -l C amqscic0.ccs
```

CICS C sample transaction

Sample C source for a CICS IBM MQ transaction is provided by `AMQSCIC0.CCS`. The transaction reads messages from the transmission queue `SYSTEM.SAMPLE.CICS.WORKQUEUE` on the

default queue manager and places them onto the local queue with a queue name that is contained in the transmission header of the message. Any failures are sent to the queue SYSTEM.SAMPLE.CICS.DLQ. Use the sample MQSC script AMQSCIC0.TST to create these queues and sample input queues.

Building your procedural application on Windows

Windows

The Windows systems publications describe how to build executable applications from the programs that you write.

This topic describes the additional tasks, and the changes to the standard tasks, that you must perform when building IBM MQ for Windows applications to run under Windows systems. ActiveX, C, C++, COBOL, and Visual Basic programming languages are supported. For information about preparing your ActiveX programs, see Using the Component Object Model Interface (WebSphere MQ Automation Classes for ActiveX). For information about preparing your C++ programs, see Using C++.

The tasks that you must perform to create an executable application using IBM MQ for Windows vary with the programming language that your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the IBM MQ for Windows include files for the language that you are using. Make yourself familiar with the contents of these files. See “IBM MQ data definition files” on page 71 for a full description.

Building 64-bit applications on Windows:

Both 32-bit and 64-bit applications are supported on IBM MQ for Windows. The IBM MQ executable and library files are supplied in both 32-bit and 64-bit forms, use the appropriate version depending on the application you are working with.

Executable files and libraries

Both 32-bit and 64-bit versions of the IBM MQ libraries are supplied in the following locations:

Table 60. Location of IBM MQ libraries

Library version	Directory containing library files
32-bit	<code>MQ_INSTALLATION_PATH \Tools\Lib</code>
64-bit	<code>MQ_INSTALLATION_PATH \Tools\Lib64</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

32-bit applications continue to work normally after migration. The 32-bit files exist in the same directory as in previous versions of the product.

If you want to create 64-bit version you must ensure that your environment is configured to use the library files in `MQ_INSTALLATION_PATH \Tools\Lib64`. Ensure that the LIB environment variable is not set to look in the folder containing the 32-bit libraries.

Preparing C programs in Windows:

Work in your typical Windows environment; IBM MQ for Windows requires nothing special.

For further information about programming 64-bit applications see Coding standards on 64-bit platforms.

- Link your programs with the appropriate libraries provided by IBM MQ:

Library file	Program/exit type
<code>MQ_INSTALLATION_PATH</code> <code>\Tools\Lib\mqm.lib</code>	server for 32-bit C
<code>MQ_INSTALLATION_PATH</code> <code>\Tools\Lib\mqic.lib</code>	client for 32-bit C
<code>MQ_INSTALLATION_PATH</code> <code>\Tools\Lib\mqicxa.lib</code>	client for 32-bit C with transaction co-ordination
<code>MQ_INSTALLATION_PATH</code> <code>\Tools\Lib64\mqm.lib</code>	server for 64-bit C
<code>MQ_INSTALLATION_PATH</code> <code>\Tools\Lib64\mqic.lib</code>	client for 64-bit C
<code>MQ_INSTALLATION_PATH</code> <code>\Tools\Lib64\mqicxa.lib</code>	client for 64-bit C with transaction co-ordination

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

The following command gives an example of compiling the sample program `amqsget0` (using the Microsoft Visual C++ compiler).

For 32-bit applications:

```
cl -MD amqsget0.c -Feamqsget.exe MQ_INSTALLATION_PATH\Tools\Lib\mqm.lib
```

For 64-bit applications:

```
cl -MD amqsget0.c -Feamqsget.exe MQ_INSTALLATION_PATH\Tools\Lib64\mqm.lib
```

Note:

- If you are writing an installable service (see the Administering for further information), you need to link to the `mqmzf.lib` library.
- If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries Encina, or BEA Tuxedo, you need to link to the `mqmxa.lib` or `mqmxa.lib` library.
- If you are writing a CICS exit, link to the `mqmcics4.lib` library.
- You must link IBM MQ libraries before any other product libraries.
- The DLLs must be in the path (PATH) that you have specified.
- If you use lowercase characters whenever possible, you can move from IBM MQ for Windows to IBM MQ on UNIX and Linux systems, where use of lowercase is necessary.

Preparing CICS and Transaction Server programs

Sample C source for a CICS IBM MQ transaction is provided by `AMQSCIC0.CCS`. You build it using the standard CICS facilities. For example, for TXSeries for Windows 2000:

1. Set the environment variable (enter the following code on one line):

```
set CICS_IBMC_FLAGS=-I MQ_INSTALLATION_PATH\Tools\C\Include;  
%CICS_IBMC_FLAGS%
```

2. Set the `USERLIB` environment variable:

```
set USERLIB=MQM.LIB;%USERLIB%
```

3. Translate, compile, and link the sample program:

```
cicstcl -l IBMC amqscic0.ccs
```

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

This is described in the *Transaction Server for Windows NT Application Programming Guide (CICS) V4*.

You can find more information about supporting CICS transactions in the Administering.

Preparing COBOL programs in Windows:

Use this information to learn to prepare COBOL programs in Windows, and preparing CICS and Transaction Server programs.

1. The 32 bit COBOL copy books are installed in the following directory: `MQ_INSTALLATION_PATH\Tools\cobol\CopyBook`.
2. The 64 bit COBOL copy books are installed in the following directory: `MQ_INSTALLATION_PATH\Tools\cobol\CopyBook64`
3. In the following examples set CopyBook to:

CopyBook

for 32 bit applications, and:

CopyBook64

for 64 bit applications.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

To prepare COBOL programs on Windows systems, link your program to one of the following libraries provided by IBM MQ:

Library file	Program or exit type
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmccb</code>	32 bit server for IBM COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcb</code>	32 bit server for Micro Focus COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqiccb</code>	32 bit client for IBM COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqiccb</code>	32 bit client for Micro Focus COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmccb</code>	64 bit server for IBM COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmcb</code>	64 bit server for Micro Focus COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqiccb</code>	64 bit client for IBM COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqiccb</code>	64 bit client for Micro Focus COBOL

When you are running a program in the MQI client environment, ensure that the DOSCALLS library appears before any COBOL or IBM MQ library.

You can use the IBM COBOL Set compiler or Micro Focus COBOL compiler depending on the program:

- Programs beginning `amqi` are suitable for the IBM COBOL Set compiler,
- Programs beginning `amqm` are suitable for the Micro Focus COBOL compiler, and
- Programs beginning `amq0` are suitable for either compiler.

IBM and Micro Focus COBOL

Relink any existing 32 bit IBM MQ Micro Focus COBOL programs using either `mqmcb.lib` or `mqiccb.lib`, rather than the `mqmccb` and `mqiccb` libraries.

To compile, for example, the sample program `amq0put0`, using IBM VisualAge COBOL:

1. Set the SYSLIB environment variable to include the path to the IBM MQ VisualAge COBOL copybooks (enter the following code on one line):

```
set SYSLIB= MQ_INSTALLATION_PATH\  
Tools\Cobol\Copybook\VAcobol;%SYSLIB%
```

2. For use on the IBM MQ server:

```
cob2 amq0put0.cbl -qlib " MQ_INSTALLATION_PATH\  
Tools\Lib\mqmcb.lib"
```

3. For use on the IBM MQ client:

```
cob2 amq0put0.cbl -qlib " MQ_INSTALLATION_PATH\  
Tools\Lib\mqiccb.lib"
```

Note: Although you must use the compiler option CALLINT(SYSTEM), this is the default for cob2.

To compile, for example, the sample program amq0put0, using Micro Focus COBOL:

1. Set the COBCPY environment variable to point to the IBM MQ COBOL copybooks (enter the following code on one line):

```
set COBCPY= MQ_INSTALLATION_PATH\  
Tools\Cobol\Copybook
```

2. Compile the program to give you an object file:

```
cob1 amq0put0 LITLINK
```

3. Link the object file to the run time system.

- Set the LIB environment variable to point to the compiler COBOL libraries.
- Link the object file for use on the IBM MQ server:

```
cbllink amq0put0.obj mqmcb.lib
```

- Or link the object file for use on the IBM MQ client:

```
cbllink amq0put0.obj mqiccb.lib
```

Preparing CICS and Transaction Server programs

To compile and link a TXSeries for Windows NT, V5.1 program using IBM VisualAge COBOL:

1. Set the environment variable (enter the following code on one line):

```
set CICS_IBMCOB_FLAGS= MQ_INSTALLATION_PATH\  
Cobol\Copybook\VAcobol;%CICS_IBMCOB_FLAGS%
```

2. Set the USERLIB environment variable:

```
set USERLIB=MQMCBB.LIB
```

3. Translate, compile, and link your program:

```
cicstcl -l IBMCOB myprog.ccp
```

This is described in the *Transaction Server for Windows NT, V4 Application Programming Guide*.

To compile and link a CICS for Windows V5 program using Micro Focus COBOL:

- Set the INCLUDE variable:

```
set  
INCLUDE=drive:\programname\ibm\websphere\tools\c\include;  
drive:\opt\cics\include;%INCLUDE%
```

- Set the COBCPY environment variable:

```
setCOBCPY=drive:\programname\ibm\websphere\tools\cobol\copybook;  
drive:\opt\cics\include
```

- Set the COBOL options:


```
- set  
- COBOPTS=/LITLINK /NOTRUNC
```


and run the following code:

```
cicstran cicsmq00.ccp
cobol cicsmq00.cbl /LITLINK /NOTRUNC
cbllink -D -Mcicsmq00 -Ocicsmq00.cbfnt cicsmq00.obj
%ICSLIB%\cicsprCBMFNT.lib user32.lib msvcrt.lib kernel32.lib mqmcb.lib
```

Preparing Visual Basic programs in Windows:

Information to consider when using Microsoft Visual Basic programs on Windows.

 From Version 9.0, IBM MQ support for Microsoft Visual Basic 6.0 is deprecated. IBM MQ classes for .NET are the recommended replacement technology. For more information, see [Developing .NET applications](#).

Note: 64-bit versions of the Visual Basic module files are not supplied.

To prepare Visual Basic programs on Windows:

1. Create a new project.
2. Add the supplied module file, CMQB.BAS, to the project.
3. Add other supplied module files if you need them:
 - CMQBB.BAS: MQAI support
 - CMQCFB.BAS: PCF support
 - CMQXB.BAS: Channel exits support
 - CMQPSB.BAS: Publish/subscribe

See “Coding in Visual Basic” on page 450 for information about using the MQCONNAny call from within Visual Basic.

Call the procedure MQ_SETDEFAULTS before making any MQI calls in the project code. This procedure sets up default structures that the MQI calls require.

Specify whether you are creating an IBM MQ server or client, before you compile or run the project, by setting the conditional compilation variable *MqType*. Set *MqType* in a Visual Basic project to 1 for a server or 2 for a client as follows:

1. Select the Project menu.
2. Select *Name Properties* (where *Name* is the name of the current project).
3. Select the Make tab in the dialog box.
4. In the Conditional Compilation Arguments field, enter this for a server:

```
MqType=1
```

or this for a client:

```
MqType=2
```

Related concepts:

“Coding in Visual Basic” on page 450

Information to consider when coding IBM MQ programs in Microsoft Visual Basic. Visual Basic is supported only on Windows.

Related reference:

“Linking Visual Basic applications with the IBM MQ MQI client code” on page 293

You can link Microsoft Visual Basic applications with the IBM MQ MQI client code on Windows.

SSPI security exit:

IBM MQ for Windows supplies a security exit for both the IBM MQ MQI client and the IBM MQ server. This is a channel-exit program that provides authentication for IBM MQ channels by using the Security Services Programming Interface (SSPI). The SSPI provides the integrated security facilities of Windows systems.

The security packages are loaded from either security.dll or secur32.dll. These DLLs are supplied with your operating system.

One-way authentication is provided using NTLM authentication services. Two-way authentication is provided using Kerberos authentication services.

The security exit program is supplied in source and object format. You can use the object code as it is, or you can use the source code as a starting point to create your own user-exit programs.

See also “Using the SSPI security exit on Windows” on page 538.

Introduction to security exits

A security exit forms a secure connection between two security exit programs, where one program is for the sending message channel agent (MCA), and one is for the receiving MCA.

The program that initiates the secure connection, that is, the first program to get control after the MCA session is established, is known as the *context initiator*. The partner program is known as the *context acceptor*.

The following table shows some of the channel types that are context initiators and their associated context acceptors.

Table 61. Context initiators and their associated context acceptors

Context Initiator	Context Acceptor
MQCHT_CLNTCONN	MQCHT_SVRCONN
MQCHT_RECEIVER	MQCHT_SENDER
MQCHT_CLUSRCVR	MQCHT_CLUSSDR

The security exit program has two entry points:

- **SCY_NTLM**

This uses NTLM authentication services, which provide one-way authentication. NTLM allows servers to verify the identities of their clients. It does not allow clients to verify a server's identity, or one server to verify the identity of another. NTLM authentication was designed for a network environment in which servers are assumed to be genuine.

- **SCY_KERBEROS**

This uses Kerberos mutual authentication services. The Kerberos protocol does not assume that servers in a network environment are genuine. Parties at both ends of a network connection can verify the identity of the other party. That is, servers can verify the identity of clients and other servers, and clients can verify the identity of a server.

What the security exit does

This topic describes what the SSPI channel-exit programs do.

The supplied channel-exit programs provide either one-way or two-way (mutual) authentication of a partner system when a session is being established. For a particular channel, each exit program has an associated *principal* (similar to a user ID, see “ IBM MQ access control and Windows principals”). A connection between two exit programs is an association between the two principals.

After the underlying session is established, a secure connection between two security exit programs (one for the sending MCA and one for the receiving MCA), is established. The sequence of operations is as follows:

1. Each program is associated with a particular principal, for example as a result of an explicit login operation.
2. The context initiator requests a secure connection with the partner from the security package (for Kerberos, the named partner) and receives a token (called token1). The token is sent, using the underlying session that is already established, to the partner program.
3. The partner program (the context acceptor) passes token1 to the security package, which verifies that the context initiator is authentic. For NTLM, the connection is now established.
4. For the Kerberos-supplied security exit (that is, for mutual authentication), the security package also generates a second token (called token2), which the context acceptor returns to the context initiator by using the underlying session.
5. The context initiator uses token2 to verify that the context acceptor is authentic.
6. At this stage, if both applications are satisfied with the authenticity of the partner's token, the secure (authenticated) connection is established.

IBM MQ access control and Windows principals

The access control that IBM MQ provides is based on the user and group. The authentication that Windows provides is based on principals, such as user and servicePrincipalName (SPN). In the case of servicePrincipalName, there might be many of these associated with a single user.

The SSPI security exit uses the relevant Windows principals for authentication. If Windows authentication is successful, the exit passes the user ID that is associated with the Windows principal to IBM MQ for access control.

The Windows principals that are relevant for authentication vary, depending on the type of authentication used.

- For NTLM authentication, the Windows principal for Context Initiator is the user ID associated with the process that is running. Because this authentication is one way, the principal associated with the Context Acceptor is irrelevant.
- For Kerberos authentication, on CLNTCONN channels, the Windows principal is the user ID associated with the process that is running. Otherwise, the Windows principal is the servicePrincipalName that is formed by adding the following prefix to the QueueManagerName.

ibmMQSeries/

Building your procedural application on z/OS

z/OS

The CICS, IMS, and z/OS publications describe how to build applications that run in these environments.

This collection of topics describes the additional tasks, and the changes to the standard tasks, that you must perform when building IBM MQ for z/OS applications for these environments. COBOL, C, C++, Assembler, and PL/I programming languages are supported. (For information about building C++ applications see Using C++.)

The tasks that you must perform to create an executable IBM MQ for z/OS application depend on both the programming language that the program is written in, and the environment in which the application will run.

In addition to coding the MQI calls in your program, add the appropriate language statements to include the IBM MQ for z/OS data definition file for the language that you are using. Make yourself familiar with the contents of these files. See “IBM MQ data definition files” on page 71 for a full description.

Note

The name **thlqual** is the high-level qualifier of the installation library on z/OS.

Preparing your program to run:

After you have written the program for your IBM MQ application to create an executable application, you have to compile or assemble it, then link-edit the resulting object code with the stub program that IBM MQ for z/OS supplies for each environment that it supports.

How you prepare your program depends on both the environment (batch, CICS, IMS(BMP or MPP), Linux or UNIX System services) in which the application runs, and the structure of the data sets on your z/OS installation.

“Dynamically calling the IBM MQ stub” on page 427 describes an alternative method of making MQI calls in your programs so that you do not need to link-edit an IBM MQ stub. This method is not available for all languages and environments.

Do not link-edit a higher level of stub program than that of the version of IBM MQ for z/OS on which your program is running. For example, a program running on MQSeries for OS/390®, V5.2 must not be link-edited with a stub program supplied with IBM MQ for z/OS V7.

Building 64 bit C applications:

In z/OS, 64 bit C applications are built using the LP64 compiler and binder options. The IBM MQ for z/OS *cmqc.h* header file recognizes when this option is provided to the compiler, and generates IBM MQ datatypes and structures appropriate for 64 bit operation.

C code built with this option must be built to use dynamic-link libraries (DLLs) appropriate for the coordination semantic required. Binding the compiled code with the appropriate side-deck defined in Side-deck name required for each coordination semantic shows the specific DLL needed.

Table 62. Side-deck name required for each coordination semantic

Coordination	Side-deck name
Single phase commit MQI	CSQBMQ2X
Two phase commit with RRS coordination, using RRS verbs	CSQBRR2X
Two phase commit with RRS coordination, using MQI verbs	CSQBRI2X

Use the EDCQCB JCL procedure, supplied with z/OS XL C/C++, to build a single phase commit IBM MQ program as a batch job, as follows:

```
//PROCS JCLLIB ORDER=CBC.SCCNPRC
//CLG EXEC EDCQCB,
// INFILE='thlqual.SCSQC37S(CSQ4BCG1)', < MQ SAMPLES
// CPARM='RENT,SSCOM,DLL,LP64,LIST,NOMAR,NOSEQ', < COMPILER OPTIONS
// LIBPRFX='CEE', < PREFIX FOR LIBRARY DSN
// LNGPRFX='CBC', < PREFIX FOR LANGUAGE DSN
// BPARM='MAP,XREF,RENT,DYNAM=DLL', < LINK EDIT OPTIONS
// OUTFILE='userid.LOAD(CSQ4BCG1),DISP=SHR'
//COMPILE.SYSLIB DD
// DD
// DD DISP=SHR,DSN=thlqual.SCSQC370
//BIND.SCSQDEFS DD DISP=SHR,DSN=thlqual.SCSQDEFS
//BIND.SYSIN DD *
INCLUDE SCSQDEFS(CSQBMQ2X)
NAME CSQ4BCG1
```

To build an RRS coordinated program in z/OS Unix System Services, compile and link as follows:

```
cc -o mqsamp -W c,LP64,DLL -W l,DYNAM=DLL,LP64 -I"//thlqual.SCSQC370" "//thlqual.SCSQDEFS(CSQBRR2X)" mqsamp.c
```

Building z/OS batch applications:

Learn how to build z/OS batch applications and the steps to consider when doing so.

To build an application for IBM MQ for z/OS that runs under z/OS batch, create job control language (JCL) that performs these tasks:

1. Compile (or assemble) the program to produce object code. The JCL for your compilation must include SYSLIB statements that make the product data definition files available to the compiler. The data definitions are supplied in the following IBM MQ for z/OS libraries:
 - For COBOL, **thlqual.SCSQCOBC**
 - For assembler language, **thlqual.SCSQMACS**
 - For C, **thlqual.SCSQC370**
 - For PL/I, **thlqual.SCSQPLIC**
2. For a C application, prelink the object code created in step 1.
3. For PL/I applications, use the compiler option EXTRN(SHORT).

- Link-edit the object code created in step 1 on page 421 (or step 2 on page 421 for a C application) to produce a load module. When you link-edit the code, you must include one of the IBM MQ for z/OS batch stub programs (CSQBSTUB or one of the RRS stub programs: CSQBRRSI or CSQBRSTB).

CSQBSTUB

single-phase commit provided by IBM MQ for z/OS

CSQBRRSI

two-phase commit provided by RRS using the MQI

CSQBRSTB

two-phase commit provided by RRS directly

Note: If you use CSQBRSTB, you must also link-edit your application with ATRSCSS from SYS1.CSSLIB. Figure 68 and Figure 69 show fragments of JCL to do this. The stubs are language-independent and are supplied in library **thlqual.SCSQLOAD**.

- Store the load module in an application load library.

```

:
:
/*
/* WEBSPHERE MQ FOR Z/OS LIBRARY CONTAINING BATCH STUB
/*
//CSQSTUB DD DSN=++THLQUAL++.SCSQLOAD,DISP=SHR
/*
:
:
//SYSIN DD *
INCLUDE CSQSTUB(CSQBSTUB)
:
:
/*

```

Figure 68. Fragments of JCL to link-edit the object module in the batch environment, using single-phase commit

```

:
:
/*
/* WEBSPHERE MQ FOR Z/OS LIBRARY CONTAINING BATCH STUB
/*
//CSQSTUB DD DSN=++THLQUAL++.SCSQLOAD,DISP=SHR
//CSSLIB DD DSN=SYS1.CSSLIB,DISP=SHR
/*
:
:
//SYSIN DD *
INCLUDE CSQSTUB(CSQBRSTB)
INCLUDE CSSLIB(ATRSCSS)
:
:
/*

```

Figure 69. Fragments of JCL to link-edit the object module in the batch environment, using two-phase commit

To run a batch or RRS program, you must include the libraries **thlqual.SCSQAUTH** and **thlqual.SCSQLOAD** in the STEPLIB or JOBLIB data set concatenation.

To run a TSO program, you must include the libraries **thlqual.SCSQAUTH** and **thlqual.SCSQLOAD** in the STEPLIB used by the TSO session.

To run a UNIX System Services batch program from the UNIX System Services shell, add the libraries **thlqual.SCSQAUTH** and **thlqual.SCSQLOAD** to the STEPLIB specification in your \$HOME?.profile like this:

```

STEPLIB= thlqual.SCSQAUTH: thlqual.SCSQLOAD
export STEPLIB

```

IBM MQ for z/OS provides a set of dynamic link libraries (DLLs) that must be used when you link-edit your applications.

There are two variants of the libraries which allow the application to use one of the following calling interfaces:

- The 31-bit Language Environment calling interface.
- The 31-bit XPLINK calling interface. z/OS XPLINK is a high performance calling convention available for C applications.

To use the DLLs, the application is bound or linked against so called *sidedecks*, instead of the stubs provided with earlier versions. The sidedecks are found in the SCSQDEFS library (instead of the SCSQLOAD library).

Table 63. Variants of dynamic link libraries

Commit	31-bit Language Environment DLL	31-bit XPLINK DLL	Equivalent stub name
1 phase commit MQI libraries	CSQBMQ1	CSQBMQ1X	CSQBSTUB
2 phase commit with RRS co-ordination using RRS transaction-control verbs	CSQBRR1	CSQBRR1X	CSQBRSTB
2 phase commit with RRS co-ordination using MQI transaction-control verbs	CSQBRI1	CSQBRI1X	CSQBRSI

Note: All sidedecks contain a definition of the data conversion entry point, MQXCNVC, previously resolved by including CSQASTUB.

Common issues:

- The following message appears on the job log if your application uses asynchronous message consume (MQCB, MQCTL or MQSUB calls) and the previous DLL interface is not used:
CSQB001E Language environment programs running in z/OS batch or USS must use the DLL interface to IBM MQ
Solution: Rebuild your application using sidedecks instead of stubs as detailed previously.
- At program build time, the following message appears
IEW2469E The Attributes of a reference to MQAPI-NAME from section your-code do not match the attributes of the target symbol
Reason: This means that you have compiled your XPLINK program with V701 (or later) version of cmqc.h, but are not binding with sidedecks.
Solution: Change your program's build file to bind against the appropriate sidedeck from SCSQDEFS instead of a stub from SCSQLOAD

The following sample JCL demonstrates how you can compile and link-edit a C program to use the 31 bit Language Environment DLL calling interface:

```
//CLG EXEC EDCCB,
// INFILE=MYPROGS.CPROGS(MYPROGRAM),
// CPARM='OPTF(DD:OPTF)',
// BPARM='XREF,MAP,DYNAM=DLL'           < LINKEDIT OPTIONS
//COMPILE.OPTF DD *
RENT,CHECKOUT(ALL),SSCOM,DEFINE(MVS),NOMARGINS,NOSEQ,DLL
SE(DD:SYSLIBV)
//COMPILE.SYSLIB DD
//                               DD
```

```

//          DD DISP=SHR,DSN=h1q.SCSQC370
//COMPILE.SYSLIBV DD DISP=SHR,DSN=h1q.BASE.H
/*
//BIND.SYSOBJ DD DISP=SHR,DSN=CEE.SCEE0BJ
//          DD DISP=SHR,DSN=h1q.SCSQDEFS
//BIND.SYSLMOD DD DISP=SHR,DSN=h1q.LOAD(MYPROGAM)
//BIND.SYSIN DD *
ENTRY CEESTART
INCLUDE SYSOBJ(CSQBMQ1)
NAME MYPROGAM(R)
//

```

Note: The compilation uses the **DLL** option. The link-edit uses **DYNAM=DLL** option and the references the **CSQBMQ1** library.

The following sample JCL demonstrates how you can compile and link-edit a C program to use the 31 bit XPLINK DLL calling interface:

```

//CLG EXEC EDCXCB,
// INFILE=MYPROGS.CPROGS(MYPROGRAM),
// CPARM='OPTF(DD:OPTF)',
// BPARM='XREF,MAP,DYNAM=DLL'          < LINKEDIT OPTIONS
//COMPILE.OPTF DD *
RENT,CHECKOUT(ALL),SSCOM,DEFINE(MVS),NOMARGINS,NOSEQ,XPLINK,DLL
SE(DD:SYSLIBV)
//COMPILE.SYSLIB DD
//          DD
//          DD DISP=SHR,DSN=h1q.SCSQC370
//COMPILE.SYSLIBV DD DISP=SHR,DSN=h1q.BASE.H
/*
//BIND.SYSOBJ DD DISP=SHR,DSN=CEE.SCEE0BJ
//          DD DISP=SHR,DSN=h1q.SCSQDEFS
//BIND.SYSLMOD DD DISP=SHR,DSN=h1q.LOAD(MYPROGAM)
//BIND.SYSIN DD *
ENTRY CEESTART
INCLUDE SYSOBJ(CSQBMQ1X)
NAME MYPROGAM(R)
//

```

Note: The compilation uses the **XPLINK** and **DLL** options. The link-edit uses **DYNAM=DLL** option and references the **CSQBMQ1X** library.

Ensure that you add the compilation option **DLL** to each program in the module. Messages such as IEW2456E 9207 SYMBOL CSQ1BAK UNRESOLVED are an indication that you need to check that all of the programs have been compiled with the **DLL** option.

Building CICS applications in z/OS:

Use this information when building CICS applications in z/OS.

To build an application for IBM MQ for z/OS that runs under CICS, you must:

- Translate the CICS commands in your program into the language in which the rest of your program is written.
- Compile or assemble the output from the translator to produce object code.
 - For PL/I programs, use the compiler option EXTRN(SHORT).
 - For C applications, if the application is not using XPLINK, use the compiler option DEFINE(MQ_OS_LINKAGE=1).
- Link-edit the object code to create a load module.

CICS provides a procedure to execute these steps in sequence for each of the programming languages it supports.

- For CICS Transaction Server for z/OS, the *CICS Transaction Server for z/OS System Definition Guide* describes how to use these procedures and the *CICS/ESA Application Programming Guide* gives more information on the translation process.

You must include:

- In the SYSLIB statement of the compilation (or assembly) stage, statements that make the product data definition files available to the compiler. The data definitions are supplied in the following IBM MQ for z/OS libraries:
 - For COBOL, **thlqual.SCSQCOBC**
 - For assembler language, **thlqual.SCSQMACS**
 - For C, **thlqual.SCSQC370**
 - For PL/I, **thlqual.SCSQPLIC**
- In your link-edit JCL, the IBM MQ for z/OS CICS stub program (CSQCSTUB). Figure 70 shows fragments of JCL code to do this. The stub is language-independent and is supplied in library **thlqual.SCSQLOAD**.

```
⋮
/*
/* WEBSPPHERE MQ FOR Z/OS LIBRARY CONTAINING CICS STUB
/*
/*CSQSTUB DD DSN=++THLQUAL++.SCSQLOAD,DISP=SHR
/*
⋮
//LKED.SYSIN DD *
INCLUDE CSQSTUB(CSQSTUB)
⋮
/*
```

Figure 70. Fragments of JCL to link-edit the object module in the CICS environment

- For CICS versions later than CICS TS 3.2, or, if you want to use IBM MQ message property APIs, or IBM MQ APIs MQCB, MQCTL, MQSTAT, MQSUB or MQSUBR, you must linkedit your object code with the CICS supplied stub, DFHMQSTB and not the IBM MQ supplied CSQCSTUB. For more information about building IBM MQ programs for CICS, see API stub program to access IBM MQ MQI calls in the CICS product documentation.

When you have completed these steps, store the load module in an application load library and define the program to CICS in the usual way.

Before you run a CICS program, your system administrator must define it to CICS as an IBM MQ program and transaction, You can then run it in the typical way.

Building IMS (BMP or MPP) applications:

Use this information when building IMS (BMP or MPP) applications.

If you are building batch DL/I programs, see “Building z/OS batch applications” on page 421. To build other applications that run under IMS (either as a BMP or an MPP), create JCL that performs these tasks:

1. Compile (or assemble) the program to produce object code. The JCL for your compilation must include SYSLIB statements that make the product data definition files available to the compiler. The data definitions are supplied in the following IBM MQ for z/OS libraries:
 - For COBOL, **th1qua1.SCSQCOBC**
 - For assembler language, **th1qua1.SCSQMACS**
 - For C, **th1qua1.SCSQC370**
 - For PL/I, **th1qua1.SCSQPLIC**
 2. For a C application, prelink the object module created in step 1.
 3. For PL/I programs, use the compiler option EXTRN(SHORT).
 4. For a C application, if the application is not using XPLINK, use the compiler option DEFINE(MQ_OS_LINKAGE=1).
 5. Link-edit the object code created in step 1 (or step 2 for a C/370 application) to produce a load module:
 - a. Include the IMS language interface module (DFSLI000).
 - b. Include the IBM MQ for z/OS IMS stub program (CSQQSTUB). Figure 71 shows fragments of JCL to do this. The stub is language independent and is supplied in library **th1qua1.SCSQLOAD**.
- Note:** If you are using COBOL, select the NODYNAM compiler option to enable the linkage editor to resolve references to CSQQSTUB unless you intend to use dynamic linking as described in “Dynamically calling the IBM MQ stub” on page 427.
6. Store the load module in an application load library.

```
⋮
/*
/* WEBSPHERE MQ FOR Z/OS LIBRARY CONTAINING IMS STUB
/*
//CSQSTUB DD DSN=th1qua1.SCSQLOAD,DISP=SHR
/*
⋮
//LKED.SYSIN DD *
INCLUDE CSQSTUB(CSQSTUB)
⋮
/*
```

Figure 71. Fragments of JCL to link-edit the object module in the IMS environment

Before you run an IMS program, your system administrator must define it to IMS as an IBM MQ program and transaction: you can then run it in the typical way.

Building z/OS UNIX System Services applications:

Use this information when building z/OS UNIX System Services applications.

To build a C application for IBM MQ for z/OS that runs under UNIX System Services, compile and link your application as follows:

```
cc -o mqsamp -W c,DLL -I "'/' thlqual.SCSQC370'" mqsamp.c "'/' thlqual.SCSQDEFS(CSQBMQ1)'"
```

where **thlqual** is the high-level qualifier used by your installation.

To run the C program, you need to add the following to your `.profile` file; this should be in your root directory:

```
STEPLIB= thlqual.SCSQANLE:thlqual.SCSQAUTH: STEPLIB
```

Note that you need to exit from UNIX System Services, and enter UNIX System Services again, for the change to be recognized.

If you want to run multiple shells, add the word `export` at the beginning of the line, that is:

```
export STEPLIB= thlqual.SCSQANLE:thlqual.SCSQAUTH: STEPLIB
```

Once this completes successfully you can link the CSQBSTUB and issue IBM MQ calls.

“Dynamically calling the IBM MQ stub” describes an alternative method of making MQI calls in your programs so that you do not need to link-edit an IBM MQ stub. This method is not available for all languages and environments.

Do not link-edit a higher level of stub program than that of the version of IBM MQ for z/OS on which your program is running. For example, a program running on , IBM WebSphere MQ for z/OS Version 7.1 must not be link-edited with a stub program supplied with IBM MQ for z/OS, Version 8.0.

Dynamically calling the IBM MQ stub:

Instead of link-editing the IBM MQ stub program with your object code, you can dynamically call the stub from within your program.

You can do this in the batch, IMS, and CICS environments. This facility is not supported in the RRS environment. If your application program uses RRS to coordinate updates, see “RRS Considerations” on page 432.

However, this method:

- Increases the complexity of your programs
- Increases the storage required by your programs at execution time
- Reduces the performance of your programs
- Means that you cannot use the same programs in other environments

If you call the stub dynamically, the appropriate stub program and its aliases must be available at execution time. To ensure this, include the IBM MQ for z/OS data set SCSQLOAD:

- For batch and IMS, in the STEPLIB concatenation of the JCL.
- For CICS, in the CICS DFHRPL concatenation.

For IMS, ensure that the library containing the dynamic stub (built as described in the information about installing the IMS adapter in Setting up the IMS adapter) is ahead of the data set SCSQLOAD in the STEPLIB concatenation of the region JCL.

Use the names shown in Table 64 when you call the stub dynamically. In PL/I, only declare the call names used in your program.

Table 64. Call names for dynamic linking

MQI call	Batch (non-RRS) dynamic call names	CICS dynamic call names	IMS dynamic call names
MQBACK	CSQBBACK	not supported	Not supported
MQBUFMH	CSQBFBMH	CSQCBFMH ¹	MQBUFMH
MQCB	CSQBCB	CSQCCB ¹	Not supported
MQCLOSE	CSQBCLOS	CSQCCLOS	MQCLOSE
MQCMIT	CSQBCOMM	not supported	Not supported
MQCONN	CSQBCONN	CSQCCONN	MQCONN
MQCONNX	CSQBCONX	CSQCCONX	MQCONNX
MQCRTMH	CSQBCTMH	CSQCCTMH ¹	MQCRTMH
MQCTL	CSQBCTL	CSQCCTL ¹	Not supported
MQDISC	CSQBDISC	CSQCDISC	MQDISC
MQDLTMH	CSQBDMH	CSQCDTMH ¹	MQDLTMH
MQDLTMP	CSQBDMH	CSQCDTMP ¹	MQDLTMP
MQGET	CSQBGET	CSQCGET	MQGET
MQINQ	CSQBINQ	CSQCINQ	MQINQ
MQINQMP	CSQBIQMP	CSQCIQMP ¹	MQINQMP
MQMHBUF	CSQBMHBF	CSQCMHBF ¹	MQMHBUF
MQOPEN	CSQBOPEN	CSQCOPEN	MQOPEN
MQPUT	CSQBPUT	CSQCPUT	MQPUT
MQPUT1	CSQBPUT1	CSQCPUT1	MQPUT1
MQSET	CSQBSET	CSQCSET	MQSET
MQSETMP	CSQBSTMP	CSQCSTMP ¹	MQSETMP
MQSTAT	CSQBSTAT	CSQCSTAT ¹	MQSTAT
MQSUB	CSQBSUB	CSQCSUB ¹	MQSUB
MQSUBRQ	CSQBSUBR	CSQCSUBR ¹	MQSUBRQ

Note: 1. These API calls are available only when using CICS TS 3.2 or later and the CSQCSTUB shipped with CICS must be used. For CICS TS 3.2, APAR PK66866 must be applied. For CICS TS 4.1, APAR PK89844 must be applied.

For examples of how to use this technique, see the following figures:

- Batch and COBOL: see Figure 72 on page 429
- CICS and COBOL: see Figure 73 on page 429
- IMS and COBOL: see Figure 74 on page 430
- Batch and assembler: see Figure 75 on page 430
- CICS and assembler: see Figure 76 on page 430
- IMS and assembler: see Figure 77 on page 430
- Batch and C: Figure 78 on page 431
- CICS and C: see Figure 79 on page 431
- IMS and C: see Figure 80 on page 431

- Batch and PL/I: see Figure 81 on page 431
- IMS and PL/I: see Figure 82 on page 432

```

...
WORKING-STORAGE SECTION.
...
    05 WS-MQOPEN                PIC X(8) VALUE 'CSQBOPEN'.
...
PROCEDURE DIVISION.
...
    CALL WS-MQOPEN WS-HCONN
                MQOD
                WS-OPTIONS
                WS-HOBJ
                WS-COMPCODE
                WS-REASON.
...

```

Figure 72. Dynamic linking using COBOL in the batch environment

```

...
WORKING-STORAGE SECTION.
...
    05 WS-MQOPEN                PIC X(8) VALUE 'CSQCOPEN'.
...
PROCEDURE DIVISION.
...
    CALL WS-MQOPEN WS-HCONN
                MQOD
                WS-OPTIONS
                WS-HOBJ
                WS-COMPCODE
                WS-REASON.
...

```

Figure 73. Dynamic linking using COBOL in the CICS environment

```

...
WORKING-STORAGE SECTION.
...
    05 WS-MQOPEN                PIC X(8) VALUE 'MQOPEN'.
...
PROCEDURE DIVISION.
...
    CALL WS-MQOPEN WS-HCONN
                MQOD
                WS-OPTIONS
                WS-HOBJ
                WS-COMPCODE
                WS-REASON.
...
* ----- *
*
*   If the compilation option 'DYNAM' is specified
*   then you may code the MQ calls as follows
*
* ----- *
...
    CALL 'MQOPEN' WS-HCONN
                MQOD
                WS-OPTIONS
                WS-HOBJ
                WS-COMPCODE
                WS-REASON.
...

```

Figure 74. Dynamic linking using COBOL in the IMS environment

```

...
LOAD EP=CSQBOPEN
...
CALL (15), (HCONN, MQOD, OPTIONS, HOBJ, COMPCODE, REASON), VL
...
DELETE EP=CSQBOPEN
...

```

Figure 75. Dynamic linking using assembly language in the batch environment

```

...
EXEC CICS LOAD PROGRAM('CSQCOPEN') ENTRY(R15)
...
CALL (15), (HCONN, MQOD, OPTIONS, HOBJ, COMPCODE, REASON), VL
...
EXEC CICS RELEASE PROGRAM('CSQCOPEN')
...

```

Figure 76. Dynamic linking using assembly language in the CICS environment

```

...
LOAD EP=MQOPEN
...
CALL (15), (HCONN, MQOD, OPTIONS, HOBJ, COMPCODE, REASON), VL
...
DELETE EP=MQOPEN
...

```

Figure 77. Dynamic linking using assembly language in the IMS environment

```

...
typedef void CALL_ME();
#pragma linkage(CALL_ME, OS)
...
main()
{
CALL_ME * csqbopen;
...
csqbopen = (CALL_ME *) fetch("CSQBOPEN");
(*csqbopen)(Hconn,&ObjDesc,Options,&Hobj,&CompCode,&Reason);
...

```

Figure 78. Dynamic linking using C language in the batch environment

```

...
typedef void CALL_ME();
#pragma linkage(CALL_ME, OS)
...
main()
{
CALL_ME * csqcopen;
...
EXEC CICS LOAD PROGRAM("CSQCOPEN") ENTRY(csqcopen);
(*csqcopen)(Hconn,&ObjDesc,Options,&Hobj,&CompCode,&Reason);
...

```

Figure 79. Dynamic linking using C language in the CICS environment

```

...
typedef void CALL_ME();
#pragma linkage(CALL_ME, OS)
...
main()
{
CALL_ME * mqopen;
...
mqopen = (CALL_ME *) fetch("MQOPEN");
(*mqopen)(Hconn,&ObjDesc,Options,&Hobj,&CompCode,&Reason);
...

```

Figure 80. Dynamic linking using C language in the IMS environment

```

...
DCL CSQBOPEN ENTRY EXT OPTIONS(ASSEMBLER INTER);
...
FETCH CSQBOPEN;

CALL CSQBOPEN(HQM,
              MQOD,
              OPTIONS,
              HOBJ,
              COMPCODE,
              REASON);

RELEASE CSQBOPEN;

```

Figure 81. Dynamic linking using PL/I in the batch environment

```

...
DCL MQOPEN  ENTRY EXT OPTIONS(ASSEMBLER INTER);
...
FETCH MQOPEN;

CALL  MQOPEN(HQM,
             MQOD,
             OPTIONS,
             HOBJ,
             COMPCODE,
             REASON);

RELEASE  MQOPEN;

```

Figure 82. Dynamic linking using PL/I in the IMS environment

RRS Considerations:

Consider using this information if your application program uses RRS to coordinate updates.

IBM MQ provides two different stubs for batch programs which need RRS coordination - see “The RRS batch adapter” on page 261. The difference in behavior of later API calls is determined at MQCONN time by the batch adapter from information passed by the stub routine on the MQCONN or MQCONNX API. This means that dynamic API calls are available for batch programs which need RRS coordination, provided that the initial connection to IBM MQ was done by using the appropriate stub. The following example illustrates this:

```

      WORKING-STORAGE SECTION.
          05 WS-MQOPEN          PIC X(8) VALUE 'MQOPEN' .
.
.
.
      PROCEDURE DIVISION.
.
.
.
      *
      * Static call to MQCONN must be resolved by linkage edit to
      * CSQBRSTB or CSQBRRSI for RRS coordination
      *
          CALL 'MQCONN' USING W00-QMGR
                          W03-HCONN
                          W03-COMPCODE
                          W03-REASON.
.
.
.
      *
          CALL WS-MQOPEN  WS-HCONN
                          MQOD
                          WS-OPTIONS
                          WS-HOBJ
                          WS-COMPCODE
                          WS-REASON.

```


Debugging your programs:

Use this information to learn about debugging TSO and CICS programs, and an insight into CICS trace.

The main aids to debugging IBM MQ for z/OS application programs are the reason codes returned by each API call. For a list of these, including ideas for corrective action, see:

- IBM MQ for z/OS messages, completion, and reason codes for IBM MQ for z/OS
- Reason codes for all other IBM MQ platforms

This topic also suggests other debugging tools to use in particular environments.

Debugging TSO programs

The following interactive debugging tools are available for TSO programs:

- TEST tool
- VS COBOL II interactive debugging tool
- INSPECT interactive debugging tool for C and PL/I programs

Debugging CICS programs

You can use the CICS Execution Diagnostic Facility (CEDF) to test your CICS programs interactively without having to modify the program or program-preparation procedure.

For more information about EDF, see the *CICS Transaction Server for z/OS CICS Application Programming Guide*.

CICS trace

You will probably also find it helpful to use the CICS Trace Control transaction (CETR) to control CICS trace activity.

For more information about CETR, see *CICS Transaction Server for z/OS CICS-Supplied Transactions* manual.

To determine whether CICS trace is active, display connection status using the CKQC panel. This panel also shows the trace number.

To interpret CICS trace entries, see Table 65.

The CICS trace entry for these values is AP0 xxx (where xxx is the trace number specified when the CICS adapter was enabled). All trace entries except CSQCTEST are issued by CSQCTRUE. CSQCTEST is issued by CSQCRST and CSQCDSF.

Table 65. CICS adapter trace entries

Name	Description	Trace sequence	Trace data
CSQCABNT	Abnormal termination	Before issuing END_THREAD ABNORMAL to IBM MQ. This is because of the end of the task and an implicit backout could be performed by the application. A ROLLBACK request is included in the END_THREAD call in this case.	Unit of work information. You can use this information when finding out about the status of work. (For example, it can be verified against the output produced by the DISPLAY THREAD command, or the IBM MQ for z/OS log print utility.)

Table 65. CICS adapter trace entries (continued)

Name	Description	Trace sequence	Trace data
CSQCBACK	Syncpoint backout	Before issuing BACKOUT to IBM MQ for z/OS. This is due to an explicit backout request from the application.	Unit of work information.
CSQCCCRC	Completion code and reason code	After unsuccessful return from API call.	Completion code and reason code.
CSQCCOMM	Syncpoint commit	Before issuing COMMIT to IBM MQ for z/OS. This can be due to a single-phase commit request or the second phase of a two-phase commit request. The request is due to an explicit syncpoint request from the application.	Unit of work information.
CSQCEXER	Execute resolve	Before issuing EXECUTE_RESOLVE to IBM MQ for z/OS.	The unit of work information of the unit of work issuing the EXECUTE_RESOLVE. This is the last indoubt unit of work in the resynchronization process.
CSQCGETW	GET wait	Before issuing CICS wait.	Address of the ECB to be waited on.
CSQCGMGD	GET message data	After successful return from MQGET.	Up to 40 bytes of the message data.
CSQCGMGH	GET message handle	Before issuing MQGET to IBM MQ for z/OS.	Object handle.
CSQCGMGI	Get message ID	After successful return from MQGET.	Message ID and correlation ID of the message.
CSQCINDL	Indoubt list	After successful return from the second INQUIRE_INDOUBT.	The indoubt units of work list.
CSQCINDO	IBM use only		
CSQCINDS	Indoubt list size	After successful return from the first INQUIRE_INDOUBT and the indoubt list is not empty.	Length of the list. Divided by 64 gives the number of indoubt units of work.
CSQCINQH	INQ handle	Before issuing MQINQ to IBM MQ for z/OS.	Object handle.
CSQCLOSH	CLOSE handle	Before issuing MQCLOSE to IBM MQ for z/OS.	Object handle.
CSQCLOST	Disposition lost	During the resynchronization process, CICS informs the adapter that it has been restarted so no disposition information regarding the unit of work being resynchronized is available.	Unit of work ID known to CICS for the unit of work being resynchronized.
CSQCNIND	Disposition not indoubt	During the resynchronization process, CICS informs the adapter that the unit of work being resynchronized should not have been indoubt (that is, perhaps it is still running).	Unit of work ID known to CICS for the unit of work being resynchronized.

Table 65. CICS adapter trace entries (continued)

Name	Description	Trace sequence	Trace data
CSQCNORT	Normal termination	Before issuing END_THREAD NORMAL to IBM MQ for z/OS. This is due to the end of the task and therefore the application might perform an implicit syncpoint commit. A COMMIT request is included in the END_THREAD call in this case.	Unit of work information.
CSQCOPNH	OPEN handle	After successful return from MQOPEN.	Object handle.
CSQCOPNO	OPEN object	Before issuing MQOPEN to IBM MQ for z/OS.	Object name.
CSQCPMGD	PUT message data	Before issuing MQPUT to IBM MQ for z/OS.	Up to 40 bytes of the message data.
CSQCPMGH	PUT message handle	Before issuing MQPUT to IBM MQ for z/OS.	Object handle.
CSQCPMGI	PUT message ID	After successful MQPUT from IBM MQ for z/OS.	Message ID and correlation ID of the message.
CSQCPREP	Syncpoint prepare	Before issuing PREPARE to IBM MQ for z/OS in the first phase of two-phase commit processing. This call can also be issued from the distributed queuing component as an API call.	Unit of work information.
CSQCP1MD	PUTONE message data	Before issuing MQPUT1 to IBM MQ for z/OS.	Up to 40 bytes of data of the message.
CSQCP1MI	PUTONE message ID	After successful return from MQPUT1.	Message ID and correlation ID of the message.
CSQCP1ON	PUTONE object name	Before issuing MQPUT1 to IBM MQ for z/OS.	Object name.
CSQCRBAK	Resolved backout	Before issuing RESOLVE_ROLLBACK to IBM MQ for z/OS.	Unit of work information.
CSQCRGMT	Resolved commit	Before issuing RESOLVE_COMMIT to IBM MQ for z/OS.	Unit of work information.
CSQCRMIR	RMI response	Before returning to the CICS RMI (resource manager interface) from a specific invocation.	Architected RMI response value. Its meaning depends of the type of the invocation. These values are documented in the <i>CICS Transaction Server for z/OS Customization Guide</i> . To determine the type of invocation, look at previous trace entries produced by the CICS RMI component.
CSQCRSYN	Resynchronization	Before the resynchronization process starts for the task.	Unit of work ID known to CICS for the unit of work being resynchronized.
CSQCSETH	SET handle	Before issuing MQSET to IBM MQ for z/OS.	Object handle.
CSQCTASE	IBM use only		

Table 65. CICS adapter trace entries (continued)

Name	Description	Trace sequence	Trace data
CSQCTEST	Trace test	Used in EXEC CICS ENTER TRACE call to verify the trace number supplied by the user or the trace status of the connection.	No data.
CSQDCFF	IBM use only		

Handling procedural program errors

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Whenever possible, the queue manager returns any errors as soon as an MQI call is made. These are *locally determined errors*.

When sending messages to a remote queue, errors might not be apparent when the MQI call is made. In this case, the queue manager that identifies the errors reports them by sending another message to the originating program. These are *remotely determined errors*.

Locally determined errors

Information about locally determined errors which include: failure on an MQI call, system interruptions, and messages containing incorrect data.

The three most common causes of errors that the queue manager can report immediately are:

- Failure of an MQI call; for example, because a queue is full
- An interruption to the running of some part of the system on which your application depends; for example, the queue manager
- Messages containing data that cannot be processed successfully

If you are using the asynchronous put facility, errors are not reported immediately. Use the MQSTAT call to retrieve status information about previous asynchronous put operations.

Failure of an MQI call

The queue manager can report immediately any errors in the coding of an MQI call. It does this using a set of predefined return codes. These are divided into completion codes and reason codes.

To show whether a call is successful, the queue manager returns a *completion code* when the call completes. There are three completion codes, indicating success, partial completion, and failure of the call. The queue manager also returns a *reason code* that indicates the reason for the partial completion or the failure of the call.

The completion and reason codes for each call are listed with the description of that call in Return codes. For more detailed information, including ideas for corrective action, see:

-  IBM MQ for z/OS messages, completion, and reason codes for IBM MQ for z/OS
- Reason codes for all other IBM MQ platforms

Design your programs to handle all the return codes that can arise from each call.

System interruptions

Your application might be unaware of any interruption if the queue manager to which it is connected has to recover from a system failure. However, you must design your application to ensure that your data is not lost if such an interruption occurs.

The methods that you can use to make sure that your data remains consistent depends on the platform on which your queue manager is running:

z/OS z/OS

In the CICS and IMS environments, you can make MQPUT and MQGET calls within units of work that are managed by CICS or IMS. In the batch environment, you can make MQPUT and MQGET calls in the same way, but you must declare sync points using:

- The IBM MQ for z/OS MQCMIT and MQBACK calls (see “Committing and backing out units of work” on page 222), or
- The z/OS Transaction Management and Recoverable Resource Manager Services (RRS) to provide two-phase sync point support. RRS allows you to update both IBM MQ and other RRS-enabled product resources, such as Db2 stored procedure resources, within a single logical unit of work. For information about RRS sync point support see “Transaction management and recoverable resource manager services” on page 226.

IBM i IBM i

You can make your MQPUT and MQGET calls within global units of work that are managed by IBM i commitment control. You can declare sync points by using the native IBM i COMMIT and ROLLBACK commands or the language-specific commands. Local units of work are managed by IBM MQ using the MQCMIT and MQBACK calls.

UNIX, Linux, and Windows systems

In these environments, you can make your MQPUT and MQGET calls in the usual way, but you must declare sync points by using the MQCMIT and MQBACK calls (see “Committing and backing out units of work” on page 222). In the CICS environment, MQCMIT and MQBACK commands are disabled, because you can make your MQPUT and MQGET calls within units of work that are managed by CICS.

Use persistent messages for carrying all data that you cannot afford to lose. Persistent messages are reinstated on queues if the queue manager has to recover from a failure. **ULW** With IBM MQ on UNIX, Linux, and Windows, an MQGET or MQPUT call within your application will fail at the point of filling all the log files, with the message MQRC_RESOURCE_PROBLEM. For more information about log files on UNIX, Linux, and Windows, see *Administering*. **z/OS** For z/OS see *Planning on z/OS*.

If the queue manager is stopped by an operator while an application is running, the quiesce option is usually used. The queue manager enters a quiescing state in which applications can continue to do work, but they must terminate as soon as convenient. Small, quick applications can probably ignore the quiescing state and continue until they terminate as normal. Longer running applications, or ones that wait for messages to arrive, should use the *fail if quiescing* option when they use the MQOPEN, MQPUT, MQPUT1, and MQGET calls. These options mean that the calls fail when the queue manager quiesces, but the application might still have time to terminate cleanly by issuing calls that ignore the quiescing state. Such applications could also commit, or back out, changes that they have made, and then terminate.

If the queue manager is forced to stop (that is, stop without quiescing), applications will receive the MQRC_CONNECTION_BROKEN reason code when they make MQI calls. Exit the application or, alternatively, on **IBM i** IBM MQ for IBM i, UNIX, Linux, and Windows systems, issue an MQDISC call.

Messages containing incorrect data

When you use units of work in your application, if a program cannot successfully process a message that it retrieves from a queue, the MQGET call is backed out.

The queue manager maintains a count (in the *BackoutCount* field of the message descriptor) of the number of times that happens. It maintains this count in the descriptor of each message that is affected. This count can provide valuable information about the efficiency of an application. Messages with backout counts that are increasing over time are being repeatedly rejected; design your application so that it analyzes the reasons for this and handles such messages accordingly.

z/OS On IBM MQ for z/OS, to make the backout count survive restarts of the queue manager, set the **HardenGetBackout** attribute to MQQA_BACKOUT_HARDENED; otherwise, if the queue manager has to restart, it does not maintain an accurate backout count for each message. Setting the attribute this way adds the penalty of extra processing.

On IBM MQ for **IBM i** IBM i, Windows, UNIX and Linux systems, the backout count always survives the queue manager restarts.

z/OS Also, on IBM MQ for z/OS, when you remove messages from a queue within a unit of work, you can mark one message so that it is not made available again if the unit of work is backed out by the application. The marked message is treated as if it has been retrieved under a new unit of work. You mark the message that is to skip backout using the MQGMO_MARK_SKIP_BACKOUT option (in the MQGMO structure) when you use the MQGET call. See “Skipping backout” on page 161 for more information about this technique.

Using report messages for problem determination

The remote queue manager cannot report errors such as failing to put a message on a queue when you make your MQI call, but it can send you a report message to say how it has processed your message.

Within your application you can create (MQPUT) report messages as well as select the option to receive them (in which case they are sent by either another application or by a queue manager).

Creating report messages

Report messages enable an application to tell another application that it cannot deal with the message that was sent.

However, the *Report* field must initially be analyzed to determine whether the application that sent the message is interested in being informed of any problems. Having determined that a report message is required, you have to decide:

- Whether you want to include the entire original message, just the first 100 bytes of data, or none of the original message.
- What to do with the original message. You can discard it or let it go to the dead-letter queue.
- Whether the contents of the *MsgId* and *CorrelId* fields are needed as well.

Use the *Feedback* field to indicate the reason for the report message being generated. Put your report messages on an application's reply-to queue. See *Feedback* for further information.

Requesting and receiving (MQGET) report messages

When you send a message to another application, you are not informed of any problems unless you complete the *Report* field to indicate the feedback that you require. See *Structure of the report field* for the options available.

Queue managers always put report messages on an application's reply-to queue and it is recommended that your own applications do the same. When you use the report message facility, specify the name of your reply-to queue in the message descriptor of your message; otherwise, the MQPUT call fails.

Your application must contain procedures that monitor your reply-to queue and process any messages that arrive on it. Remember that a report message can contain all the original message, the first 100 bytes of the original message, or none of the original message.

The queue manager sets the *Feedback* field of the report message to indicate the reason for the error; for example, the target queue does not exist. Your programs should do the same.

For more information about report messages, see “Report messages” on page 11.

Remotely determined errors

When you send messages to a remote queue, even when the local queue manager has processed your MQI call without finding an error, other factors can influence how your message is handled by a remote queue manager.

For example, the queue that you are targeting might be full, or might not even exist. If your message has to be handled by other intermediate queue managers on the route to the target queue, any of these could find an error.

Problems delivering a message

When an MQPUT call fails, you can try to put the message on the queue again, return it to the sender, or put it on the dead-letter queue.

Each option has its merits, but you might not want to try putting a message again if the reason that the MQPUT failed was because the destination queue was full. In this instance, putting it on the dead-letter queue allows you to deliver it to the correct destination queue later on.

Retry message delivery

Before the message is put on a dead-letter queue, a remote queue manager attempts to put the message on the queue again if the attributes *MsgRetryCount* and *MsgRetryInterval* have been set for the channel, or if there is a retry exit program for it to use (the name of which is held in the channel attribute *MsgRetryExitId* field).

If the *MsgRetryExitId* field is blank, the values in the attributes *MsgRetryCount* and *MsgRetryInterval* are used.

If the *MsgRetryExitId* field is not blank, the exit program of this name runs. For more information about using your own exit programs, see “Channel-exit programs for messaging channels” on page 341.

Return message to sender

You return a message to the sender by requesting a report message to be generated to include all of the original message.

See “Report messages” on page 11 for details on report message options.

Using the dead-letter (undelivered message) queue:

When a queue manager cannot deliver a message, it attempts to put the message on its dead-letter queue. This queue should be defined when the queue manager is installed.

Your programs can use the dead-letter queue in the same way that the queue manager uses it. You can find the name of the dead-letter queue by opening the queue manager object (using the MQOPEN call) and inquiring about the **DeadLetterQName** attribute (using the MQINQ call).

When the queue manager puts a message on this queue, it adds a header to the message, the format of which is described by the dead-letter header (MQDLH) structure; see MQDLH - Dead-letter header. This header includes the name of the target queue and the reason that the message was put on the dead-letter queue. It must be removed and the problem must be resolved before the message is put on the intended queue. Also, the queue manager changes the *Format* field of the message descriptor (MQMD) to indicate that the message contains an MQDLH structure.

MQDLH structure

You are recommended to add an MQDLH structure to all messages that you put on the dead-letter queue; however, if you intend to use the dead-letter handler provided by certain IBM MQ products, you must add an MQDLH structure to your messages.

The addition of the header to a message might make the message too long for the dead-letter queue, so always make sure that your messages are shorter than the maximum size allowed for the dead-letter queue, by at least the value of the MQ_MSG_HEADER_LENGTH constant. The maximum size of messages allowed on a queue is determined by the value of the **MaxMsgLength** attribute of the queue. For the dead-letter queue, make sure that this attribute is set to the maximum allowed by the queue manager. If your application cannot deliver a message, and the message is too long to be put on the dead-letter queue, follow the advice given in the description of the MQDLH structure.

Ensure that the dead-letter queue is monitored, and that any messages arriving on it get processed. The dead-letter queue handler runs as a batch utility and can be used to perform various actions on selected messages on the dead-letter queue. For further details, see "Dead-letter queue processing" on page 441.

If data conversion is necessary, the queue manager converts the header information when you use the MQGMO_CONVERT option on the MQGET call. If the process putting the message is an MCA, the header is followed by all the text of the original message.

Messages put on the dead-letter queue might be truncated if they are too long for this queue. A possible indication of this situation is the messages on the dead-letter queue being the same length as the value of the **MaxMsgLength** attribute of the queue.

Dead-letter queue processing:

This information contains general-use programming interface information when using dead-letter queue processing.

Dead-letter queue processing depends on local system requirements, but consider the following things when you draw up the specification:

- The message can be identified as having a dead-letter queue header because the value of the format field in the MQMD, is MQFMT_DEAD_LETTER_HEADER.
- On IBM MQ for z/OS using CICS, if an MCA puts this message to the dead-letter queue, the *PutApplType* field is MQAT_CICS, and the *PutApplName* field is the *ApplId* of the CICS system followed by the transaction name of the MCA.
- The reason for the message to be routed to the dead-letter queue is contained in the *Reason* field of the dead-letter queue header.
- The dead-letter queue header contains details of the destination queue name and queue manager name.
- The dead-letter queue header contains fields that have to be reinstated in the message descriptor before the message is put to the destination queue. These are:
 1. *Encoding*
 2. *CodedCharSetId*
 3. *Format*
- The message descriptor is the same as PUT by the original application, except for the three fields shown (Encoding, CodedCharSetId, and Format).

Your dead-letter queue application must do one or more of the following things:

- Examine the *Reason* field. A message might have been put by an MCA for the following reasons:
 - The message was longer than the maximum message size for the channel
The reason is MQRC_MSG_TOO_BIG_FOR_CHANNEL
 - The message could not be put to its destination queue
The reason is any MQRC_* reason code that can be returned by an MQPUT operation
 - A user exit has requested this action
The reason code is that supplied by the user exit, or the default MQRC_SUPPRESSED_BY_EXIT
- Try to forward the message to its intended destination, where this is possible.
- Retain the message for a certain length of time before discarding when the reason for the diversion is determined, but not immediately correctable.
- Give instructions to administrators correct problems where these have been determined.
- Discard messages that are corrupted or otherwise not processible.

There are two ways to deal with the messages that you have recovered from the dead-letter queue:

1. If the message is for a local queue:
 - Carry out any code translations required to extract the application data
 - Carry out code conversions on that data if this is a local function
 - Put the resulting message on the local queue with all the detail of the message descriptor restored
2. If the message is for a remote queue, put the message on the queue.

For information about how undelivered messages are handled in a distributed queuing environment, see What happens when a message cannot be delivered?.

Multicast programming

Use this information to learn about the IBM MQ Multicast programming tasks such as connecting to a queue manager and exception reporting.

IBM MQ Multicast was designed to be as transparent to the user as possible and yet still be compatible with existing applications. Defining a COMMINFO object and setting the TOPIC object's **MCAST** and **COMMINFO** parameters, means that existing IBM MQ applications do not require substantial rewriting to use multicast. However, there might be some limitations (see “Multicast and the MQI” for more information) and some security issues to consider (see Multicast security for more information).

Multicast and the MQI

Use this information to understand the major Message Queue Interface (MQI) concepts and how they relate to IBM MQ Multicast.

Multicast subscriptions are nondurable; because there are no physical queues involved, there is nowhere to store the offline messages that are created by durable subscriptions.

After an application has subscribed to a multicast topic, it is given back an object handle which it can consume or MQGET from, as if it were a handle to a queue. This means that only managed multicast subscriptions (subscriptions created with MQSO_MANAGED) are supported, that is; it is not possible to make a subscription and 'point' the messages at a queue. This means that messages must be consumed from the object handle returned on the subscription call. On the client, the messages are stored in a message buffer until they are consumed by the client; see MessageBuffer stanza of the client configuration file for more information. If the client does not keep up with the publishing rate, the messages are discarded as required, with the oldest messages discarded first.

It is normally an administration decision whether an application uses Multicast or not, specified by setting the MCAST attribute of a TOPIC object. If a publishing application must ensure that multicast is not used, it can use the MQOO_NO_MULTICAST option. Similarly, a subscribing application can ensure that multicast is not used by subscribing with the MQSO_NO_MULTICAST option.

IBM MQ Multicast supports the use of message selectors. A selector is used by an application to register its interest in only those messages with properties that satisfy the SQL92 query that the selection string represents. For more information about message selectors, see “Selectors” on page 22.

The following table lists all the major MQI concepts and how they relate to Multicast:

Table 66. MQI concepts and how they relate to multicast

MQI Concept	Action when tried using multicast	Reason code
Putting a zero length message	Rejected	2005 (07D5) (RC2005): MQRC_BUFFER_LENGTH_ERROR
Grouping	Rejected	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR
Segmentation	Rejected	2443 (098B) (RC2443): MQRC_SEGMENTATION_NOT_ALLOWED
Distribution lists	Rejected	2154 (086A) (RC2154): MQRC_RECS_PRESENT_ERROR
MQINQ	Rejected for topics handles: MQINQ and MQSET of topics is not supported.	2038 (07F6) (RC2038): MQRC_NOT_OPEN_FOR_INQUIRE

Table 66. MQI concepts and how they relate to multicast (continued)

MQI Concept	Action when tried using multicast	Reason code
MQINQ	Accepted for managed handle. Only Current Depth can be inquired.	<ul style="list-style-type: none"> If the value is Current Depth, then there is no applicable reason code. If the value is anything other than Current Depth, the reason code is 2067 (0813) (RC2067): MQRC_SELECTOR_ERROR.
MQSET	Rejected for all handles.	2040 (07F8) (RC2040): MQRC_NOT_OPEN_FOR_SET
Transactions (XA or not)	Rejected	2072 (0818) (RC2072): MQRC_SYNCPOINT_NOT_AVAILABLE
Message browse	Rejected	2036 (07F4) (RC2036): MQRC_NOT_OPEN_FOR_BROWSE
Lock messages	Rejected	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR
Browse with mark	Rejected	2036 (07F4) (RC2036): MQRC_NOT_OPEN_FOR_BROWSE
Pass context	Rejected	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR
MQPUT1	Rejected. It is invalid to try and MQPUT1 to a Multicast only topic.	2560 (0A00) (RC2560): MQRC_MULTICAST_ONLY
Durable subscription	Rejected if the topic is marked as "Multicast only", otherwise a non-Multicast subscription is made.	2436 (0984) (RC2436): MQRC_DURABILITY_NOT_ALLOWED
TopicString > 255	Rejected. If the topic string is greater than 255 characters, it is rejected in the client.	2425 (0979) (RC2425): MQRC_TOPIC_STRING_ERROR
Non-managed subscription made	Rejected if the topic is marked as "Multicast only", otherwise a non-Multicast subscription is made.	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR
MQPMO_NOT_OWN_SUBS	Rejected	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR

The following items expand on some of the MQI concepts from the previous table, and provides information on some of the MQI concepts that are not in the table:

Message persistence

For nondurable multicast subscribers, persistent messages from the publisher are delivered in an unrecoverable fashion.

Message truncation

Message truncation is supported, which means that it is possible for an application to:

1. Issue an MQGET.

2. Get MQRC_TRUNCATED_MSG_FAILED.
3. Allocate a larger buffer.
4. Reissue the MQGET to retrieve the message.

Subscription expiry

Subscription expiry is not supported. Any attempt to set an expiry is ignored.

High availability for multicast

Use this information to understand IBM MQ Multicast continuous peer-to-peer operation; although IBM MQ connects to an IBM MQ queue manager, messages do not flow through that queue manager.

Although a connection to a queue manager must be made in order to MQOPEN or MQSUB the multicast topic object, the messages themselves do not flow through the queue manager. Therefore, after the MQOPEN or MQSUB is completed on the multicast topic object, it is possible to continue transmitting multicast messages even if the connection to the queue manager has been lost. There are two modes of operation:

A normal connection is made to the queue manager

Multicast communication is possible while the connection to the queue manager exists. If the connection fails, the normal MQI rules are applied, for example; an MQPUT to the multicast object handle returns 2009 (07D9) (RC2009): MQRC_CONNECTION_BROKEN.

A reconnecting client connection is made to the queue manager

Multicast communication is possible even during the reconnection cycle. This means that even when the connection to the queue manager has been broken, the putting and consuming of multicast messages is not affected. The client attempts to reconnect to a queue manager, and if that reconnection fails, the connection handle becomes broken and all MQI calls, including multicast ones, fail. For more information, see: Automatic client reconnection

If any application explicitly issues an MQDISC, then all multicast subscriptions and object handles are closed.

Multicast continuous peer-to-peer operation

One of the advantages of peer-to-peer communication between the clients is that the messages do not need to flow through the queue manager; therefore if the connection to the queue manager breaks, message transfer continues. The following restrictions apply to the continuous message requirements of this mode:

- The connection must be made using one of the MQCNO_RECONNECT_* options for continuous operation. This process means that although the communications session might be broken, the actual connection handle is not broken, and is in the reconnecting state instead. If reconnection fails, the connection handle is now broken which prevents all further MQI calls.
- Only MQPUT, MQGET, MQINQ, and Async Consume are supported in this mode. Any MQOPEN, MQCLOSE, or MQDISC verbs require reconnection to the queue manager to complete.
- Status flows to the queue manager stop; any state in the queue manager might therefore be stale or missing. This means that the clients might be sending and receiving messages and there is no status known on the queue manager. For more information, see: Multicast application monitoring

Data conversion in the MQI for multicast messaging

Use this information to understand how data conversion works for IBM MQ Multicast messaging.

IBM MQ Multicast is a shared, connectionless protocol, and so it is not possible for each client to make specific requests for data conversion. Every client subscribed to the same multicast stream receives the same binary data; therefore, if IBM MQ data conversion is required, the conversion is performed locally at each client.

Data is converted on the client for IBM MQ Multicast traffic. If the **MQGMO_CONVERT** option is specified, data conversion is done as requested. User defined formats need the data conversion exit installed on the client; see "Writing data-conversion exits" on page 364 for information about which libraries are now in the client and server packages.

For information about administering data conversion, see [Enabling data conversion for Multicast messaging](#).

For more information about data conversion, see [Data conversion](#).

For more information about data conversion exits and ClientExitPath, see ClientExitPath stanza of the client configuration file.

Multicast exception reporting

Use this information to learn about IBM MQ Multicast event handlers and reporting IBM MQ Multicast exceptions.

IBM MQ Multicast assists with problem determination by calling the event handler to report multicast events which are reported using the standard IBM MQ event handler mechanism.

An individual Multicast event can result in more than one IBM MQ event being called because there might be multiple MQHCONN connection handles using the same multicast transmitter or receiver. However, each multicast exception causes only one event handler to be called per IBM MQ connection.

The IBM MQ MQCBDO_EVENT_CALL constant enables applications to register a callback to receive only IBM MQ events, and the MQCBDO_MC_EVENT_CALL enable applications to register a callback to receive only multicast events. If both constants are used, both types of event are received.

Requesting Multicast events

IBM MQ Multicast events use the MQCBDO_MC_EVENT_CALL constant in the `cbd.Options` field. The following example demonstrates how to request multicast events:

```
cbd.CallbackType      = MQCBT_EVENT_HANDLER;
cbd.Options           = MQCBDO_MC_EVENT_CALL;
cbd.CallbackFunction = EventHandler;
MQCB(Hcon,MQOP_REGISTER,&cbd,MQHO_UNUSABLE_HOBJ,NULL,NULL,&CompCode,&Reason);
```

When the MQCBDO_MC_EVENT_CALL option is specified for the `cbd.Options` field, the event handler is sent only IBM MQ Multicast events instead of connection level events. To request that both types of events are sent to the event handler, the application must specify the MQCBDO_EVENT_CALL constant in the `cbd.Options` field as well as the MQCBDO_MC_EVENT_CALL constant as shown in the following example:

```
cbd.CallbackType      = MQCBT_EVENT_HANDLER;
cbd.Options           = MQCBDO_EVENT_CALL | MQCBDO_MC_EVENT_CALL
cbd.CallbackFunction = EventHandler;
MQCB(Hcon,MQOP_REGISTER,&cbd,MQHO_UNUSABLE_HOBJ,NULL,NULL,&CompCode,&Reason);
```

If neither of these constants is used, only connection level events are sent to the event handler.

For more information about values for the Options field see Options (MQLONG).

Multicast event format

IBM MQ Multicast exceptions include some supporting information which is returned in the **Buffer** parameter of the callback function. The **Buffer** pointer points to an array of pointers and the MQCBC.DataLength field specifies the size, in bytes, of the array. The first element of the array always points to a short text description of the event. More parameters might be supplied depending on the type of event. The following table lists the exceptions:

Table 67. Multicast event code descriptions

Event code	Description	Additional data
MQMCEV_PACKET_LOSS	Unrecoverable packet loss	Number of lost packets
MQMCEV_HEARTBEAT_TIMEOUT	Long absence of heartbeat control packet	N/A
MQMCEV_VERSION_CONFLICT	Reception of newer protocol version packets	N/A
MQMCEV_RELIABILITY	Different reliability modes of the transmitter and the receiver	N/A
MQMCEV_CLOSED_TRANS	Topic transmission is closed by 1 source	N/A
MQMCEV_STREAM_ERROR	Error detected on stream	N/A
MQMCEV_NEW_SOURCE	A new source starts to transmit on the topic	Source structure
MQMCEV_RECEIVE_QUEUE_TRIMMED	Packets removed from PacketQ due to time or space expiration	Number of trimmed packets
MQMCEV_PACKET_LOSS_NACK_EXPIRE	Unrecoverable packet loss due to NACK expiration	Number of lost packets
MQMCEV_ACK_RETRIES_EXCEEDED	Packets removed from history after max_ack_retries was exceeded	Number of packets removed
MQMCEV_STREAM_SUSPEND_NACK	NACKs have been suspended on a stream accepted by this topic	Suspend stream ID Time in milliseconds that the stream is suspended for
MQMCEV_STREAM_RESUME_NACK	NACKs have been resumed after they have been suspended on a stream	Stream ID
MQMCEV_STREAM_EXPELLED	A stream accepted by this topic has been rejected due to an expel request	Stream ID
MQMCEV_FIRST_MESSAGE	First message from a source	Message number
MQMCEV_LATE_JOIN_FAILURE	Failed to start late join session	N/A
MQMCEV_MESSAGE_LOSS	Unrecoverable message loss	Number of lost messages
MQMCEV_SEND_PACKET_FAILURE	Multicast transmitter failed to send a multicast packet	N/A
MQMCEV_REPAIR_DELAY	Multicast receiver did not receive a repair packet for an outstanding NAK	N/A
MQMCEV_MEMORY_ALERT_ON	Receiver reception buffers are filling up	Buffer pool utilization percentage

Table 67. Multicast event code descriptions (continued)

Event code	Description	Additional data
MQMCEV_MEMORY_ALERT_OFF	Receiver reception buffers are down to normal	Buffer pool utilization percentage
MQMCEV_NACK_ALERT_ON	Receiver repair packet request rate reached high water mark	Current repair request rate in packets per second
MQMCEV_NACK_ALERT_OFF	Receiver repair packet request rate is down to normal	Current repair request rate in packets per second
MQMCEV_REPAIR_ALERT_ON	Transmitter repair packet send rate reached high water mark	N/A
MQMCEV_REPAIR_ALERT_OFF	Transmitter repair packet send rate is down to normal	N/A
MQMCEV_SHM_DEST_UNUSABLE	The Shared Memory region used by a transmitter topic destination has been detected to be unusable	N/A
MQMCEV_SHM_PORT_UNUSABLE	The Shared Memory port used by a receiver instance has been detected to be unusable	N/A
MQMCEV_CCT_GETTIME_FAILED	The get time from Coordinated Cluster Time failed	N/A
MQMCEV_DEST_INTERFACE_FAILURE	The network interface used by a transmitter topic destination has failed and a backup network interface is unavailable	
MQMCEV_DEST_INTERFACE_FAILOVER	The network interface used by a transmitter topic destination has failed and a successful failover to another Interface has been completed	
MQMCEV_PORT_INTERFACE-FAILURE	The network interface used by a receiver rmmPort has failed and a backup network interface is unavailable (or has also failed)	RMM configuration
MQMCEV_PORT_INTERFACE_FAILOVER	The network interface used by a receiver rmmPort has failed and a successful failover to another Interface has been completed	RMM configuration

Coding in C

Note the information in the following sections when coding IBM MQ programs in C.

- “Parameters of the MQI calls”
- “Parameters with undefined data type”
- “Data types”
- “Manipulating binary strings” on page 449
- “Manipulating character strings” on page 449
- “Initial values for structures” on page 449
- “Initial values for dynamic structures” on page 450
- “Use from C++” on page 450

Parameters of the MQI calls

Parameters that are *input-only* and of type MQHCONN, MQHOBJ, MQHMSG, or MQLONG are passed by value; for all other parameters, the *address* of the parameter is passed by value.

Not all parameters that are passed by address need to be specified every time a function is invoked. Where a particular parameter is not required, a null pointer can be specified as the parameter on the function invocation, in place of the address of the parameter data. Parameters for which this is possible are identified in the call descriptions.

No parameter is returned as the value of the function; in C terminology, this means that all functions return void.

The attributes of the function are defined by the MQENTRY macro variable; the value of this macro variable depends on the environment.

Parameters with undefined data type

The MQGET, MQPUT, and MQPUT1 functions each have a **Buffer** parameter that has an undefined data type. This parameter is used to send and receive the application's message data.

Parameters of this sort are shown in the C examples as arrays of MQBYTE. You can declare the parameters in this way, but it is typically more convenient to declare them as the structure that describes the layout of the data in the message. The function parameter is declared as a pointer-to-void, and so the address of any data can be specified as the parameter on the function invocation.

Data types

All data types are defined with the typedef statement.

For each data type, the corresponding pointer data type is also defined. The name of the pointer data type is the name of the elementary or structure data type prefixed with the letter P to denote a pointer. The attributes of the pointer are defined by the MQPOINTER macro variable; the value of this macro variable depends on the environment. The following code illustrates how to declare pointer data types:

```
#define MQPOINTER          /* depends on environment */
...
typedef MQLONG  MQPOINTER PMQLONG; /* pointer to MQLONG */
typedef MQMD   MQPOINTER PMQMD;   /* pointer to MQMD */
```


Manipulating binary strings

Strings of binary data are declared as one of the MQBYTE n data types.

Whenever you copy, compare, or set fields of this type, use the C functions `memcpy`, `memcmp`, or `memset`:

```
#include <string.h>
#include "cmqc.h"

MQMD MyMsgDesc;

memcpy(MyMsgDesc.MsgId,          /* set "MsgId" field to nulls */
       MQMI_NONE,              /* ...using named constant */
       sizeof(MyMsgDesc.MsgId));

memset(MyMsgDesc.CorrelId,      /* set "CorrelId" field to nulls */
       0x00,                   /* ...using a different method */
       sizeof(MQBYTE24));
```

Do not use the string functions `strcpy`, `strcmp`, `strncpy`, or `strncmp` because these do not work correctly with data declared as MQBYTE24.

Manipulating character strings

When the queue manager returns character data to the application, the queue manager always pads the character data with blanks to the defined length of the field. The queue manager does not return null-terminated strings, but you can use them in your input. Therefore, when copying, comparing, or concatenating such strings, use the string functions `strncpy`, `strncmp`, or `strncat`.

Do not use the string functions that require the string to be terminated by a null (`strcpy`, `strcmp`, and `strcat`). Also, do not use the function `strlen` to determine the length of the string; use instead the `sizeof` function to determine the length of the field.

Initial values for structures

The include file `<cmqc.h>` defines various macro variables that you can use to provide initial values for the structures when declaring instances of those structures. These macro variables have names of the form `MQxxx_DEFAULT`, where `MQxxx` represents the name of the structure. Use them like this:

```
MQMD MyMsgDesc = {MQMD_DEFAULT};
MQPMO MyPutOpts = {MQPMO_DEFAULT};
```

For some character fields, the MQI defines particular values that are valid (for example, for the *StrucId* fields or for the *Format* field in MQMD). For each of the valid values, two macro variables are provided:

- One macro variable defines the value as a string with a length, excluding the implied null, that exactly matches the defined length of the field. For example, the symbol $\hat{\Delta}$ represents a blank character:

```
#define MQMD_STRUC_ID "MD $\hat{\Delta}$ "
#define MQFMT_STRING "MQSTR $\hat{\Delta}$  $\hat{\Delta}$ "
```

Use this form with the `memcpy` and `memcmp` functions.

- The other macro variable defines the value as an array of `char`; the name of this macro variable is the name of the string form suffixed with `_ARRAY`. For example:

```
#define MQMD_STRUC_ID_ARRAY 'M','D',' $\hat{\Delta}$ ',' $\hat{\Delta}$ '
#define MQFMT_STRING_ARRAY 'M','Q','S','T','R',' $\hat{\Delta}$ ',' $\hat{\Delta}$ ',' $\hat{\Delta}$ '
```

Use this form to initialize the field when an instance of the structure is declared with values different from those provided by the `MQMD_DEFAULT` macro variable.

Initial values for dynamic structures

When a variable number of instances of a structure are required, the instances are typically created in main storage obtained dynamically using the `calloc` or `malloc` functions.

To initialize the fields in such structures, the following technique is recommended:

1. Declare an instance of the structure using the appropriate `MQxxx_DEFAULT` macro variable to initialize the structure. This instance becomes the *model* for other instances:

```
MQMD ModelMsgDesc = {MQMD_DEFAULT};  
/* declare model instance */
```

Code the `static` or `auto` keywords on the declaration to give the model instance static or dynamic lifetime, as required.

2. Use the `calloc` or `malloc` functions to obtain storage for a dynamic instance of the structure:

```
PMQMD InstancePtr;  
InstancePtr = malloc(sizeof(MQMD));  
/* get storage for dynamic instance */
```

3. Use the `memcpy` function to copy the model instance to the dynamic instance:

```
memcpy(InstancePtr,&ModelMsgDesc,sizeof(MQMD));  
/* initialize dynamic instance */
```

Use from C++

For the C++ programming language, the header files contain the following additional statements that are included only when a C++ compiler is used:

```
#ifdef __cplusplus  
extern "C" {  
#endif  
  
/* rest of header file */  
  
#ifdef __cplusplus  
}  
#endif
```

Coding in Visual Basic

Windows

Information to consider when coding IBM MQ programs in Microsoft Visual Basic. Visual Basic is supported only on Windows.

Note: From IBM WebSphere MQ Version 7.0, outside the .NET environment, support for Visual Basic (VB) has been stabilized at the Version 6.0 level. Most new function added to IBM WebSphere MQ Version 7.0 or later is not available to VB applications. If you are programming in VB.NET, use the IBM MQ classes for .NET. For more information, see *Developing .NET applications*.

V 9.0.0

From Version 9.0, IBM MQ support for Microsoft Visual Basic 6.0 is deprecated. IBM MQ classes for .NET are the recommended replacement technology.

To avoid unintended translation of binary data passing between Visual Basic and IBM MQ, use an `MQBYTE` definition instead of `MQSTRING`. `CMQB.BAS` defines several new `MQBYTE` types that are equivalent to a C byte definition and uses these within IBM MQ structures. For example, for the `MQMD` (message descriptor) structure, `MsgId` (message identifier) is defined as `MQBYTE24`.

Visual Basic does not have a pointer data type, so references to other IBM MQ data structures are by offset rather than pointer. Declare a compound structure consisting of the two component structures, and specify the compound structure on the call. IBM MQ support for Visual Basic provides an MQCONNAny call to make this possible and allow client applications to specify the channel properties on a client connection. It accepts an untyped structure (MQCNOCD) in place of the typical MQCNO structure.

The MQCNOCD structure is a compound structure consisting of an MQCNO followed by an MQCD. This structure is declared in the exits header file CMQXB. Use the routine MQCNOCD_DEFAULTS to initialize an MQCNOCD structure. A sample making MQCONN calls is provided (amqscnxb.vbp).

MQCONNAny has the same parameters as MQCONN, except that the **ConnectOpts** parameter is declared as being of Any data type rather than of MQCNO data type. This allows the function to accept either the MQCNO or the MQCNOCD structure. This function is declared in the main header file CMQB.

Related concepts:

“Preparing Visual Basic programs in Windows” on page 417

Information to consider when using Microsoft Visual Basic programs on Windows.

Related reference:

“Linking Visual Basic applications with the IBM MQ MQI client code” on page 293

You can link Microsoft Visual Basic applications with the IBM MQ MQI client code on Windows.

Coding in COBOL

Note the information in the following section when coding IBM MQ programs in COBOL.

Named constants

The names of constants are shown containing the underscore character (`_`) as part of the name. In COBOL, you must use the hyphen character (`-`) in place of the underscore. Constants that have character-string values use the single quotation mark character (`'`) as the string delimiter. To make the compiler accept this character, use the compiler option APOST.

The copy file CMQV contains declarations of the named constants as level-10 items. To use the constants, declare the level-01 item explicitly, then use the COPY statement to copy in the declarations of the constants:

```
WORKING-STORAGE SECTION.  
01 MQM-CONSTANTS.  
COPY CMQV.
```

However, this method causes the constants to occupy storage in the program even if they are not referred to. If the constants are included in many separate programs within the same run unit, multiple copies of the constants will exist; this might result in a significant amount of main storage being used. You can avoid this by adding the GLOBAL clause to the level-01 declaration:

```
* Declare a global structure to hold the constants  
01 MQM-CONSTANTS GLOBAL.  
COPY CMQV.
```

This allocates storage for only *one* set of constants within the run unit; the constants, however, can be referred to by *any* program within the run unit, not just the program that contains the level-01 declaration.

Ensuring structure alignment

Care should be taken to ensure IBM MQ structures that are passed on to start on the MQ call's must be aligned on word boundaries. A word boundary is 4 bytes for 32-bit processes, 8 bytes for 64 bit processes and 16 bytes for 128 bit processes (IBM i).

Where possible, place all IBM MQ structures together so they are all boundary aligned.

Coding in System/390 assembler language (Message queue interface)

Note the information in the following sections when coding IBM MQ for z/OS programs in assembler language.

- “Names”
- “Using the MQI calls”
- “Declaring constants”
- “Specifying the name of a structure” on page 453
- “Specifying the form of a structure” on page 453
- “Controlling the listing” on page 453
- “Specifying initial values for fields” on page 454
- “Writing reenterable programs” on page 454
- “Using CEDF” on page 454

Names

The names of parameters in the descriptions of calls, and the names of fields in the descriptions of structures are shown in mixed case. In the assembler-language macros supplied with IBM MQ, all names are in uppercase.

Using the MQI calls

The MQI is a call interface, so assembler-language programs must observe the OS linkage convention.

In particular, before they issue an MQI call, assembler-language programs must point register R13 at a save area of at least 18 full words. This save area provides storage for the called program. It stores the registers of the caller before their contents are destroyed, and restores the contents of the caller's registers on return.

Note: This is important for CICS assembler-language programs that use the DFHEIENT macro to set up their dynamic storage, but that choose to override the default DATAREG from R13 to other registers. When the CICS Resource Manager Interface receives control from the stub, it saves the current contents of the registers at the address to which R13 is pointing. Failing to reserve a save area for this purpose gives unpredictable results, and will probably cause an abend in CICS.

Declaring constants

Most constants are declared as equates in macro CMQA.

However, the following constants cannot be defined as equates, and these are not included when you call the macro using default options:

- MQACT_NONE
- MQCI_NONE
- MQFMT_NONE
- MQFMT_ADMIN
- MQFMT_COMMAND_1
- MQFMT_COMMAND_2
- MQFMT_DEAD_LETTER_HEADER
- MQFMT_EVENT
- MQFMT_IMS

- MQFMT_IMS_VAR_STRING
- MQFMT_PCF
- MQFMT_STRING
- MQFMT_TRIGGER
- MQFMT_XMIT_Q_HEADER
- MQMI_NONE

To include them, add the keyword EQUONLY=NO when you call the macro.

CMQA is protected against multiple declaration, so you can include it many times. However, the keyword EQUONLY takes effect only the first time that the macro is included.

Specifying the name of a structure

To allow more than one instance of a structure to be declared, the macro that generates the structure prefixes the name of each field with a user-specifiable string and an underscore character (_).

Specify the string when you invoke the macro. If you do not specify a string, the macro uses the name of the structure to construct the prefix:

```
* Declare two object descriptors
CMQODA      Prefix used="MQOD_" (the default)
MY_MQOD CMQODA      Prefix used="MY_MQOD_"
```

The structure declarations in Call descriptions show the default prefix.

Specifying the form of a structure

The macros can generate structure declarations in one of two forms, controlled by the DSECT parameter:

DSECT=YES

An assembler-language DSECT instruction is used to start a new data section; the structure definition immediately follows the DSECT statement. No storage is allocated, so no initialization is possible. The label on the macro invocation is used as the name of the data section; if no label is specified, the name of the structure is used.

DSECT=NO

Assembler-language DC instructions are used to define the structure at the current position in the routine. The fields are initialized with values, which you can specify by coding the relevant parameters on the macro invocation. Fields for which no values are specified on the macro invocation are initialized with default values.

DSECT=NO is assumed if the DSECT parameter is not specified.

Controlling the listing

You can control the appearance of the structure declaration in the assembler-language listing with the LIST parameter:

LIST=YES

The structure declaration appears in the assembler-language listing.

LIST=NO

The structure declaration does not appear in the assembler-language listing. This is assumed if the LIST parameter is not specified.

Specifying initial values for fields

You can specify the value to be used to initialize a field in a structure by coding the name of that field (without the prefix) as a parameter on the macro invocation, accompanied by the value required.

For example, to declare a message descriptor structure with the *MsgType* field initialized with MQMT_REQUEST, and the *ReplyToQ* field initialized with the string MY_REPLY_TO_QUEUE, use the following code:

```
MY_MQMD    CMQMDA    MSGTYPE=MQMT_REQUEST,      X
REPLYTOQ=MY_REPLY_TO_QUEUE
```

If you specify a named constant (or equate) as a value on the macro invocation, use the CMQA macro to define the named constant. You must not enclose in single quotation marks (' ') values that are character strings.

Writing reenterable programs

IBM MQ uses its structures for both input and output. If you want your program to remain reenterable:

1. Define working storage versions of the structures as DSECTs, or define the structures inline within an already-defined DSECT. Then copy the DSECT to storage that is obtained using:
 - For batch and TSO programs, the STORAGE or GETMAIN z/OS assembler macros
 - For CICS, the working storage DSECT (DFHEISTG) or the EXEC CICS GETMAIN command

To correctly initialize these working storage structures, copy a constant version of the corresponding structure to the working storage version.

Note: The MQMD and MQXQH structures are each more than 256 bytes long. To copy these structures to storage, use the MVCL assembler instruction.

2. Reserve space in storage by using the LIST form (MF=L) of the CALL macro. When you use the CALL macro to make an MQI call, use the EXECUTE form (MF=E) of the macro, using the storage reserved earlier, as shown in the example under “Using CEDF.” For more examples of how to do this, see the assembler language sample programs as shipped with IBM MQ.

Use the assembler language RENT option to help you to determine if your program is reenterable.

For information on writing reenterable programs, see *MVS/ESA Application Development Guide: Assembler Language Programs*, GC28-1644.

Using CEDF

If you want to use the CICS-supplied transaction, CEDF (CICS Execution Diagnostic Facility) to help you to debug your program, add the ,VL keyword to each CALL statement, for example:

```
CALL MQCONN, (NAME,HCONN,COMPCODE,REASON),MF=(E,PARMAREA),VL
```

The previous example is reenterable assembler-language code where PARMAREA is an area in the working storage that you specified.

Using the MQI calls

The MQI is a call interface, so assembler-language programs must observe the OS linkage convention. In particular, before they issue an MQI call, assembler-language programs must point register R13 at a save area of at least 18 full words. This save area provides storage for the called program. It stores the registers of the caller before their contents are destroyed, and restores the contents of the caller's registers on return.

Note: This is important for CICS assembler-language programs that use the DFHEIENT macro to set up their dynamic storage, but that choose to override the default DATAREG from R13 to other registers. When the CICS Resource Manager Interface receives control from the stub, it saves the current contents of the registers at the address to which R13 is pointing. Failing to reserve a proper save area for this purpose gives unpredictable results, and will probably cause an abend in CICS.

Coding IBM MQ programs in RPG (IBM i only)



In the IBM MQ documentation, the parameters of calls, the names of data types, the fields of structures, and the names of constants are all described using their long names. In RPG, these names are abbreviated to six or fewer uppercase characters.

For example, the field *MsgType* becomes *MDMT* in RPG. For more information, see the IBM i Application Programming Reference (ILE/RPG).

Coding in pTAL

Note the information in the following section when coding IBM MQ programs in pTAL.

HP Integrity NonStop Server

Defining and initializing IBM MQ structures

pTAL structure definitions for IBM MQ structures are provided with names that end with ^DEF. For example, the following pTAL declarations would be coded to create a IBM MQ Message Descriptor (MQMD) structure and a IBM MQ Put Message Options (MQPMO) structure.

```
STRUCT MYMD(MQMD^DEF);    ! Declare an MQMD structure
STRUCT MYPMO(MQPMO^DEF); ! Declare an MQPMO structure
```

IBM MQ provides pTAL DEFINE with names that end with ^DEFAULT to initialize IBM MQ structures with default values. The following pTAL statements are coded to assign default values to the declared MQMD and the MQPMO structures:

```
MQMD^DEFAULT(MYMD);      ! Assign default values to an MQMD structure
MQPMO^DEFAULT(MYPMO);    ! Assign default values to an MQPMO structure
```

You can declare and initialize other IBM MQ structures by using similar code.

pTAL and the CRE

pTAL programs cannot initialize the Common Runtime Environment and therefore they must be used with a C-language or COBOL main routine.

The pTAL samples that are provided with IBM MQ use a C-language mainline routine that is called AMQSPTM0.C

Parameters with MQCHAR data type

The MQGET, MQPUT, and MQPUT1 procedures each have a **Buffer** parameter that has an MQCHAR .EXT data type. This parameter is used to send and receive the application's message data.

Parameters of this sort are shown in the pTAL samples as arrays of string. You can declare the parameters in this way, but it is typically more convenient to declare them as the structure that describes the layout of the data in the message. The procedure parameter is declared as an MQCHAR .EXT, but the address of any data can be specified as the parameter on the procedure invocation.

Manipulating character strings

When the queue manager returns character data to the application, the queue manager always pads the character data with blanks to the defined length of the field. The queue manager does not return null-terminated strings, but you can use them in your input.

Coding in PL/I (z/OS only)

Useful information when coding for IBM MQ in PL/I.

Structures

Structures are declared with the BASED attribute, and so do not occupy any storage unless the program declares one or more instances of a structure.

An instance of a structure can be declared using the like attribute, for example:

```
dc1 my_mqmd    like MQMD; /* one instance */
dc1 my_other_mqmd like MQMD; /* another one */
```

The structure fields are declared with the INITIAL attribute; when the like attribute is used to declare an instance of a structure, that instance inherits the initial values defined for that structure. You need to set only those fields where the value required is different from the initial value.

PL/I is not sensitive to case, and so the names of calls, structure fields, and constants can be coded in lowercase, uppercase, or mixed case.

Named constants

The named constants are declared as macro variables; as a result, named constants that are not referred to by the program do not occupy any storage in the compiled procedure.

However, the compiler option that causes the source to be processed by the macro preprocessor must be specified when the program is compiled.

All the macro variables are character variables, even the ones that represent numeric values. Although this might seem counter intuitive, it does not result in any data-type conflict after the macro variables have been substituted by the macro processor, for example:

```
%dc1 MQMD_STRUC_ID char;
%MQMD_STRUC_ID = 'MD';
```

```
%dc1 MQMD_VERSION_1 char;
%MQMD_VERSION_1 = '1';
```


Using the IBM MQ sample procedural programs

These sample programs are written in procedural languages, and demonstrate typical uses of the Message Queue Interface (MQI). IBM MQ programs on different platforms.



About this task

There are two sets of samples:

- Sample programs for distributed systems and IBM i.
- Sample programs for z/OS.

Procedure

Use the following links to find out more about the sample programs:

-  “Using the sample programs on Multiplatforms”
-  “Using the sample programs for z/OS” on page 568

Related concepts:

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

“Developing applications for IBM MQ” on page 1

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM MQ support applications written in procedural languages, and object oriented languages and frameworks.

“Design considerations for IBM MQ applications” on page 43

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

“Writing a procedural application for queuing” on page 74

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

“Writing client procedural applications” on page 285

What you need to know to write client applications on IBM MQ using a procedural language.

“Writing publish/subscribe applications” on page 172

Start writing publish/subscribe IBM MQ applications.

“Building a procedural application” on page 382

You can write an IBM MQ application in one of several procedural languages, and run the application on several different platforms.

“Handling procedural program errors” on page 436

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Related tasks:

“Developing web services with IBM MQ” on page 1217

You can develop IBM MQ applications for web services using the IBM MQ transport for SOAP.

Using the sample programs on Multiplatforms



These sample procedural programs are delivered with the product. The samples are written in C and COBOL, and demonstrate typical uses of the Message Queue Interface (MQI).

About this task

The samples are not intended to demonstrate general programming techniques, so some error checking that you might want to include in a production program is omitted.

The source code for all the samples is provided with the product; this source includes comments that explain the message queuing techniques demonstrated in the programs.

 For RPG programming, see IBM i Application Programming Reference (ILE/RPG).

The names of the samples start with the prefix amq. The fourth character indicates the programming language, and the compiler where necessary:

- s: C language
- 0: COBOL language on both IBM and Micro Focus compilers
- i: COBOL language on IBM compilers only
- m: COBOL language on Micro Focus compilers only

The eighth character of the executable indicates whether the sample runs in local binding mode or client mode. If there is no eighth character, then the sample runs in local bindings mode. If the eighth character is 'c' then the sample runs in client mode.

Before you can run the sample applications, you must first create and configure a queue manager. To set up the queue manager to accept client connections, see “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470.

Procedure

Use the following links to find out more about the sample programs:

- “Features demonstrated in the sample programs on Multiplatforms” on page 459
- “The Publish/Subscribe sample programs” on page 498
- “The Put sample programs” on page 503
- “The Distribution List sample program” on page 490
- “The Browse sample programs” on page 478
- “The Browser sample program” on page 480
- “The Get sample programs” on page 492
- “The Reference Message sample programs” on page 505
- “The Request sample programs” on page 512
- “The Inquire sample programs” on page 496
- “The Inquire Properties of a Message Handle sample program” on page 497
- “The Set sample programs” on page 518
- “The Echo sample programs” on page 491
- “The Data-Conversion sample program” on page 482
- “The Triggering sample programs” on page 522
- “The Asynchronous Put sample program” on page 477
- “Database coordination samples” on page 483
- “The CICS transaction sample” on page 481
- “Using the TUXEDO samples on UNIX and Windows” on page 524
- “Dead-letter queue handler sample” on page 489
- “The Connect sample program” on page 481
- “The API exit sample program” on page 475

- “Using the SSPI security exit on Windows” on page 538
- “Running the samples using remote queues” on page 539
- “The Cluster Queue Monitoring sample program (AMQSCLM)” on page 540
- “Sample program for Connection Endpoint Lookup (CEPL)” on page 550

Related concepts:

“C++ sample programs” on page 798

Four sample programs are supplied, to demonstrate getting and putting messages.

Related tasks:

► **z/OS** “Using the sample programs for z/OS” on page 568

The sample procedural applications that are delivered with IBM MQ for z/OS demonstrate typical uses of the Message Queue Interface (MQI).

Features demonstrated in the sample programs on Multiplatforms: **Multi**

A collection of tables that show the techniques demonstrated by the IBM MQ sample programs.

All the samples open and close queues using the MQOPEN and MQCLOSE calls, so these techniques are not listed separately in the tables. See the heading that includes the platform that you are interested in.

► **z/OS** For the z/OS platform, see “Using the sample programs for z/OS” on page 568.

Samples for UNIX and Linux systems: **UNIX** **Linux**

The techniques demonstrated by the sample programs for IBM MQ on UNIX and Linux systems.

See “Preparing and running sample programs on UNIX and Linux” on page 473 to find out where the sample programs for IBM MQ on UNIX and Linux systems are stored.

Table 68 The table lists which C and COBOL source files are provided, and whether a server or client executable is included.

Table 68. IBM MQ on UNIX and Linux sample programs demonstrating use of the MQI (C and COBOL)

Technique	C (source) (1 on page 461)	COBOL (source) (2 on page 461)	Server (C executable)	Client (C executable) (3 on page 461)
Using the publish/subscribe interface	amqspuba amqssuba amqssbxa	no sample	amqspub amqssub amqssbx	no sample
Putting messages using the MQPUT call	amqsput0	amq0put0	amqsput	amqsputc
Putting a single message using the MQPUT1 call	amqsinqa amqsecha	amqminqx amqmechx amqiinqx amqiechx	amqsinq amqsech	amqsechc
Putting messages to a distribution list (4 on page 461)	amqsptl0	amq0ptl0.cbl	amqsptl	amqsptlc
Replying to a request message	amqsinqa	amqminqx amqiinqx	amqsinq	no sample
Getting messages using browse (no wait)	amqsgbr0	amq0gbr0	amqsgbr	no sample
Getting messages (wait with a time limit)	amqsget0	amq0get0	amqsget	amqsgetc
Getting messages (unlimited wait)	amqstrg0	no sample	amqstrg	amqstrgc

Table 68. IBM MQ on UNIX and Linux sample programs demonstrating use of the MQI (C and COBOL) (continued)

Technique	C (source) (1 on page 461)	COBOL (source) (2 on page 461)	Server (C executable)	Client (C executable) (3 on page 461)
Getting messages (with data conversion)	amqsecha	no sample	amqsech	no sample
Putting Reference Messages to a queue (4 on page 461)	amqsprma	no sample	amqsprm	amqsprmc
Getting Reference Messages from a queue (4 on page 461)	amqsgrma	no sample	amqsgrm	amqsgrmc
Reference Message channel exit (4 on page 461)	amqsqrma amqsxrma	no sample	amqsxrm	no sample
Browsing first 20 characters of a message	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Browsing complete messages	amqsbcg0	no sample	amqsbcg	amqsbcgc
Using a shared input queue	amqsinqa	amqminqx amqiinqx	amqsinq	amqsinqc
Using an exclusive input queue	amqstrg0	amq0req0	amqstrg	amqstrgc
Using the MQINQ call	amqsinqa	amqminqx amqiinqx	amqsinq	no sample
Using the MQSET call	amqsseta	amqmsetx amqisetx	amqsset	amqssetc
Using a reply-to queue	amqsreq0	amq0req0	amqsreq	amqsreqc
Requesting message exceptions	amqsreq0	amq0req0	amqsreq	no sample
Accepting a truncated message	amqsgbr0	amq0gbr0	amqsgbr	no sample
Using a resolved queue name	amqsgbr0	amq0gbr0	amqsgbr	no sample
Triggering a process	amqstrg0	no sample	amqstrg	amqstrgc
Using data conversion	(5 on page 461)	no sample	no sample	no sample
IBM MQ (coordinating XA-compliant database managers) accessing a single database using SQL	amqsxas0.sqc Db2 amqsxas0.ec Informix	amq0xas0.sqb	no sample	no sample
IBM MQ (coordinating XA-compliant database managers) accessing two databases using SQL	amqsxag0.c amqsxab0.sqc amqsxaf0.sqc	amq0xag0.cbl amq0xab0.sqb amq0xaf0.sqb	no sample	no sample
CICS transaction (6 on page 461)	amqscic0.ccs	no sample	amqscic0	no sample
Encina transaction (4 on page 461)	amqsxae0	no sample	amqsxae0	no sample
TUXEDO transaction to put messages (7 on page 461)	amqstpx	no sample	no sample	no sample
TUXEDO transaction to get messages (7 on page 461)	amqstgx	no sample	no sample	no sample
Server for TUXEDO (7 on page 461)	amqstxsx	no sample	no sample	no sample
Dead-letter queue handler	Directory ./tools/c/ Samples/dlq (8 on page 461)	no sample	amqsdld	no sample
From an MQI client, putting a message	no sample	no sample	no sample	amqsputc
From an MQI client, getting a message	no sample	no sample	no sample	amqsgetc

Table 68. IBM MQ on UNIX and Linux sample programs demonstrating use of the MQI (C and COBOL) (continued)

Technique	C (source) (1)	COBOL (source) (2)	Server (C executable)	Client (C executable) (3)
Connecting to the queue manager using MQCONN	amqscnxc	no sample	no sample	amqscnxc
Using API exits	amqsaxe0	no sample	amqsaxe	no sample
Cluster workload balancing exit	amqswlm0	no sample	amqswlm	no sample
Putting messages asynchronously and getting status using the MQSTAT call	amqsapt0	no sample	amqsapt	amqsaptc
Reconnectable clients	amqsphac amqsghac amqsmhac	no sample	not applicable	amqsphac amqsghac amqsmhac
Using message consumers to asynchronously consume messages from multiple queues	amqscbf0	no sample	amqscbf	amqscbfc
Specifying TLS connection information on MQCONN	amqssslc	no sample	not applicable	amqssslc

Notes:

1. The executable version of the IBM MQ MQI client samples share the same source as the samples that run in a server environment.
2. Compile programs beginning 'amqm' with the Micro Focus COBOL compiler, those beginning 'amqi' with the IBM COBOL compiler, and those beginning 'amq0' with either.
3. The executable versions of the IBM MQ MQI client samples are not available on IBM MQ for HP-UX.
4. Supported on IBM MQ for AIX, IBM MQ for HP-UX, and IBM MQ for Solaris only.
5. On IBM MQ for AIX, IBM MQ for HP-UX, and IBM MQ for Solaris this program is called amqsvfc0.c
6. CICS is supported by IBM MQ for AIX and IBM MQ for HP-UX only.
7. TUXEDO is not supported by IBM MQ for Linux on System p.
8. The source for the dead-letter queue handler consists of several files and is provided in a separate directory.

For detailed information about support for UNIX and Linux systems, see System Requirements for IBM MQ.

Samples for IBM MQ client for HP Integrity NonStop Server: **NSS Client**

The techniques demonstrated by the sample programs for the IBM MQ client on HP Integrity NonStop Server systems.

Table 69 lists which C, COBOL, and pTAL source sample programs are provided.

Table 69. IBM MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL

Technique	C				COBOL		pTAL	
	OSS (Source)	OSS (Executable)	Guardian (Source)	Guardian (Executable)	OSS (Source)	Guardian (Source)	OSS (Source)	Guardian (Source)
Using the publish/subscribe interface	amqspuba.c amqssbxa.c amqssuba.c amqspse0.c	amqspubc amqssbxc amqssubc	MQSPUBC MQSSBXC MQSSUBC	AMQSPUBC AMQSSBXC AMQSSUBC	amq0pub0.c amq0sub0.c	MQSPUBL MQSSUBL	amqtpub0.ta amqtsub0.ta	MQSPUBT MQSSUBT

Table 69. IBM MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)

Technique	C				COBOL		pTAL	
Putting messages using the MQPUT call	amqsput0.c	amqsputc	MQSPUTC	AMQSPUTC	amq0put0.cb	MQSPUTL	amqtput0.tal	MQSPUTT
Putting a single message using the MQPUT1 call	amqsecha.c	amqsechc	MQSECHC	AMQSECHC			amqtech0.tal	MQSECHT
Putting messages to a distribution list	amqsptl0.c	amqsptlc	MQSPTLC	AMQSPTLC	amq0ptl0.cb	MQSPTLL		
Replying to a request message	amqsinqa.c	amqsinqc	MQSINQC	AMQSINQC				
Getting messages (no wait)	amqsgbr0.c	amqsgbrc	MQSGBRC	AMQSGBRC	amq0gbr0.cb	MQSGBRL		
Getting messages (wait with a time limit)	amqsget0.c	amqsgetc	MQSGETC	AMQSGETC	amq0get0.cb	MQSGETL	amqtget0.tal	MQSGETT
Getting messages (unlimited wait)	amqstrg0.c	amqstrgc	MQSTRGC	AMQSTRGC				
Getting messages (with data conversion)	amqsecha.c	amqsechc	MQSECHC	AMQSECHC				
Putting Reference Messages to a queue	amqsprma.c	amqsprmc	MQSPRMC	AMQSPRMC				
Getting Reference Messages from a queue	amqsgrma.c	amqsgrmc	MQSGRMC	AMQSGRMC				
Reference Message channel exit	amqsqrma.c amqsxrma.c		MQSQRMC MQSXRMC					

Table 69. IBM MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)


Technique	C				COBOL		pTAL	
Browsing first 20 characters of a message	amqsgbr0.c	amqsgbrc	MQSGBRC	AMQSGBRC	amq0gbr0.cb	MQSGBRL		
Browsing complete messages	amqsbcg0.c	amqsbcgc	MQSBCGC	AMQSBCGC				
Using a shared input queue	amqsinqa.c	amqsinqc	MQSINQC	AMQSINQC				
Using an exclusive input queue	amqstrg0.c	amqstrgc	MQSTRGC	AMQSTRGC				
Using the MQINQ call	amqsinqa.c	amqsinqc	MQSINQC	AMQSINQC				
Using the MQSET call	amqsseta.c	amqssetc	MQSSETC	AMQSSETC				
Using a reply-to queue	amqsreq0.c	amqsreqc	MQSREQC	AMQSREQC	amq0req0.cb	MQSREQL		
Requesting message exceptions	amqsreq0.c	amqsreqc	MQSREQC	AMQSREQC	amq0req0.cb	MQSREQL		
Accepting a truncated message	amqsgbr0.c	amqsgbrc	MQSGBRC	AMQSGBRC	amq0gbr0.cb	MQSGBRL		
Using a resolved queue name	amqsgbr0.c	amqsgbrc	MQSGBRC	AMQSGBRC	amq0gbr0.cb	MQSGBRL		
Triggering a process	amqstrg0.c	amqstrgc	MQSTRGC	AMQSTRGC				
Using data conversion	amqsvfc0.c							
Dead letter queue handler (1)	Directory ./samp/dlq							
Connecting to a queue manager using MQCONN	amqscnxc.c	amqscnxc	MQSCNXC					

Table 69. IBM MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)

Technique	C				COBOL		pTAL	
Using API exits	amqsaxe0.c amqsaem0.c							
Cluster workload balancing exit	amqswlm0.c		MQSWLMC					
Cluster queue monitor	amqsclma.c							
Putting messages asynchronously and getting status using the MQSTAT call	amqsapt0.c	amqsaptc	MQSAPTC	MQSAPTC				
Reconnecting clients	amqsgnac.c amqsmnac.c amqspnac.c	amqsgnac amqsmnac amqspnac	MQSGHAC MQSMHAC MQSPHAC MQSFHAC	AMQSGHAC AMQSMHAC AMQSPHAC AMQSFHAC				
Using the message consumers to asynchronously consume messages from multiple queues	amqscbf0.c	amqscbfc						
Specifying TLS connection information on MQCONN	amqssslc.c	amqssslc	MQSSSLC	AMQSSSLC				
Activity trace	amqsact0.c	amqsactc	MQSACTC	AMQSACTC				
Message properties	amqsiqma.c amqsstma.c	amqsiqmc amqsstmc	MQSIQMC MQSSTMC	AMQSIQMC AMQSSTMC				
Command server	amqsstop.c		MQSSTOC					
Log events	amqslog0.c	amqslogc	MQSLOGC	AMQSLOGC				
Accounting	amqsmon0.c	amqsmonc	MQSMONC	AMQSMONC				
Administrative interface	amqsaicq.c amqsaiem.c amqsailq.c							

Table 69. IBM MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)

Technique	C				COBOL		pTAL	
An example of a C language main function for invoking pTAL			MQSPTMC					
Notes:								
1. The source for the dead-letter queue handler consists of several files and is provided in a separate directory.								
2. For information about developing applications for your IBM MQ client on the HP Integrity NonStop Server platform, see:								
<ul style="list-style-type: none"> • “Building your procedural application on HP Integrity NonStop Server” on page 388 <ul style="list-style-type: none"> – “Preparing C programs in HP Integrity NonStop Server” on page 390 – “Preparing COBOL programs for the IBM MQ client for HP Integrity NonStop Server” on page 391 – “Preparing pTAL programs” on page 392 								

Samples for IBM MQ for Windows: 

The techniques demonstrated by the sample programs for IBM MQ for Windows.

Table 70 lists which C and COBOL source files are provided, and whether a server or client executable is included.

Table 70. IBM MQ for Windows sample programs demonstrating use of the MQI (C and COBOL)

Technique	C (source)	COBOL (source)	Server (C executable)	Client (C executable)
Using the publish/subscribe interface	amqspuba amqssuba amqssbxa	no sample	amqspub amqssub amqssbx	no sample
Putting messages using the MQPUT call	amqsput0	amq0put0	amqsput	amqsputc
Putting a single message using the MQPUT1 call	amqsinqa amqsecha	amqminq2 amqmech2 amqiinq2 amqiech2	amqsinq amqsech	amqsinqc amqsechc
Putting messages to a distribution list	amqsptl0	amq0ptl0.cbl	amqsptl	amqsptlc
Replying to a request message	amqsinqa	amqminq2 amqiinq2	amqsinq	amqsinqc
Getting messages (no wait)	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Getting messages (wait with a time limit)	amqsget0	amq0get0	amqsget	amqsgetc
Getting messages (unlimited wait)	amqstrg0	no sample	amqstrg	amqstrgc
Getting messages (with data conversion)	amqsecha	no sample	amqsech	amqsechc
Putting Reference Messages to a queue	amqsprma	no sample	amqsprm	amqsprmc
Getting Reference Messages from a queue	amqsgrma	no sample	amqsgrm	amqsgrmc
Reference Message channel exit	amqsqrma amqsxrma	no sample	amqsxrm	no sample

Table 70. IBM MQ for Windows sample programs demonstrating use of the MQI (C and COBOL) (continued)

Technique	C (source)	COBOL (source)	Server (C executable)	Client (C executable)
Browsing first 20 characters of a message	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Browsing complete messages	amqsbcg0	no sample	amqsbcg	amqsbcgc
Using a shared input queue	amqsinqa	amqminq2 amqiinq2	amqsinq	amqsinqc
Using an exclusive input queue	amqstrg0	amq0req0	amqstrg	amqstrgc
Using the MQINQ call	amqsinqa	amqminq2 amqiinq2	amqsinq	amqsinqc
Using the MQSET call	amqsseta	amqmset2 amqiset2	amqsset	amqssetc
Using the MQINQMP call	amqsiqma	no sample	no sample	no sample
Using a reply-to queue	amqsreq0	amq0req0	amqsreq	amqsreqc
Requesting message exceptions	amqsreq0	amq0req0	amqsreq	amqsreqc
Accepting a truncated message	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Using a resolved queue name	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Triggering a process	amqstrg0	no sample	amqstrg	amqstrgc
Using data conversion	amqsvfc0	no sample	no sample	no sample
IBM MQ (coordinating XA-compliant database managers) accessing a single database using SQL	amqsxas0.sqc Db2 amqsxas0.ec Informix	amq0xas0.sqb	no sample	no sample
IBM MQ (coordinating XA-compliant database managers) accessing two databases using SQL	amqsxag0.c amqsxab0.sqc Db2 amqsxaf0.sqc Db2	amq0xag0.cbl amq0xab0.sqb amq0xaf0.sqb	no sample	no sample
TUXEDO transaction to put messages	amqstpx	no sample	no sample	no sample
TUXEDO transaction to get messages	amqstgx	no sample	no sample	no sample
Server for TUXEDO	amqstxsx	no sample	no sample	no sample
Dead-letter queue handler	Directory ./tools/c/ Samples/dlq (1 on page 467)	no sample	amqsdlq	no sample
From an IBM MQ MQI client, putting a message	no sample	no sample	no sample	amqsputc
From an IBM MQ MQI client, getting a message	no sample	no sample	no sample	amqsgetc
Connecting to the queue manager using MQCONN	amqscnxc	no sample	no sample	amqscnxc
Using API exits	amqsaxe0	no sample	amqsaxe	no sample
Cluster workload balancing	amqswlm0	no sample	amqswlm	no sample
SSPI security routines	amqsspin	no sample	amqrspin.dll	amqrspin.dll
Putting messages asynchronously and getting status using the MQSTAT call	amqsapt0	no sample	amqsapt	amqsaptc

Table 70. IBM MQ for Windows sample programs demonstrating use of the MQI (C and COBOL) (continued)

Technique	C (source)	COBOL (source)	Server (C executable)	Client (C executable)
Reconnectable clients	amqsphac amqsgnac amqsmnac	no sample	Not applicable	amqsphac amqsgnac amqsmnac
Using message consumers to asynchronously consume messages from multiple queues	amqscbf0	no sample	amqscbf	amqscbfc
Specifying TLS connection information on MQCONN	amqssslc	no sample	not applicable	amqssslc

Notes:

1. The source for the dead-letter queue handler consists of several files and is provided in a separate directory.

Visual Basic samples for IBM MQ for Windows: 

The techniques demonstrated by the sample programs for IBM MQ on Windows systems.

Table 71 shows the techniques demonstrated by the IBM MQ for Windows sample programs.

A project can contain several files. When you open a project within Visual Basic, the other files are loaded automatically. No executable programs are provided.

All the sample projects, except mqtrivc.vbp, are set up to work with the IBM MQ server. To find out how to change the sample projects to work with the IBM MQ clients see “Preparing Visual Basic programs in Windows” on page 417.

Table 71. IBM MQ for Windows sample programs demonstrating use of the MQI (Visual Basic)

Technique	Project file name
Putting messages using the MQPUT call	amqsputb.vbp
Getting messages using the MQGET call	amqsgetb.vbp
Browsing a queue using the MQGET call	amqsbcb.vbp
Simple MQGET and MQPUT sample (client)	mqtrivc.vbp
Simple MQGET and MQPUT sample (server)	mqtrivb.vbp
Putting and getting strings and user-defined structures using MQPUT and MQGET	strings.vbp
Using PCF structures to start and stop a channel	pcfsamp.vbp
Creating a queue using the MQAI	amqsaicq.vbp
Listing a queue manager's queues using the MQAI	amqsailq.vbp
Monitoring events using the MQAI	amqsaie.vbp

The techniques demonstrated by the sample programs for IBM MQ on IBM i systems.

Table 72 shows the techniques demonstrated by the IBM MQ for IBM i sample programs. Some techniques occur in more than one sample program, but only one program is listed in the table.

Table 72. Sample programs demonstrating use of the MQI (C and COBOL) on IBM i

Technique	C (source) (1 on page 469)	COBOL (source) (2 on page 469)	RPG (source) (3 on page 469)	Client (C executable)(4)
Putting messages using the MQPUT call	AMQSPUT0	AMQ0PUT4	AMQ3PUT4	AMQSPUTC
Putting messages from a data file using the MQPUT call	AMQSPUT4	no sample	no sample	no sample
Putting a single message using the MQPUT1 call	AMQSINQ4, AMQSECH4	AMQ0INQ4, AMQ0ECH4	AMQ3INQ4, AMQ3ECH4	AMQSINQC, AMQSECHC
Putting messages to a distribution list	AMQSPTL4	no sample	no sample	AMQSPTLC
Replying to a request message	AMQSINQ4	AMQ0INQ4	AMQ3INQ4	AMQSINQC
Getting messages (no wait)	AMQSGBR4	AMQ0GBR4	AMQ3GBR4	AMQSGBRC
Getting messages (wait with a time limit)	AMQSGET4	AMQ0GET4	AMQ3GET4	AMQSGETC
Getting messages (unlimited wait)	AMQSTRG4	no sample	AMQ3TRG4	AMQSTRGC
Getting messages (with data conversion)	AMQSECH4	AMQ0ECH4	AMQ3ECH4	AMQSECHC
Putting Reference Messages to a queue	AMQSPRM4	no sample	no sample	AMQSPRMC
Getting Reference Messages from a queue	AMQSGRM4	no sample	no sample	AMQSGRMC
Reference Message channel exit	AMQSORM4, AMQSXRM4	no sample	no sample	no Sample
Message exit	AMQSCMX4	no sample	no sample	no Sample
Browsing first 49 characters of a message	AMQSGBR4	AMQ0GBR4	AMQ3GBR4	AMQSGBRC
Browsing complete messages	AMQSBCG4	no sample	no sample	AMQSBCGC
Using a shared input queue	AMQSINQ4	AMQ0INQ4	AMQ3INQ4	AMQSINQC
Using an exclusive input queue	AMQSREQ4	AMQ0REQ4	AMQ3REQ4	AMQSREQC
Using the MQINQ call	AMQSINQ4	AMQ0INQ4	AMQ3INQ4	AMQSINQC
Using the MQSET call	AMQSSET4	AMQ0SET4	AMQ3SET4	AMQSSETC
Using a reply-to queue	AMQSREQ4	AMQ0REQ4	AMQ3REQ4	AMQSREQC
Requesting message exceptions	AMQSREQ4	AMQ0REQ4	AMQ3REQ4	AMQSREQC
Accepting a truncated message	AMQSGBR4	AMQ0GBR4	AMQ3GBR4	AMQSGBRC
Using a resolved queue name	AMQSGBR4	AMQ0GBR4	AMQ3GBR4	AMQSGBRC
Triggering a process	AMQSTRG4	no sample	AMQ3TRG4	AMQSTRGC
Trigger server	AMQSERV4	no sample	AMQ3SRV4	no sample
Using a trigger server (including CICS transactions)	AMQSERV4	no sample	AMQ3SRV4	no sample
Using data conversion	AMQSVFC4	no sample	no sample	no sample
Using API exits	AMQSAXE0	no sample	no sample	no sample
Cluster workload balancing	AMQSWLM0	no sample	no sample	no sample
Putting messages asynchronously and getting status using the MQSTAT call	AMQSAPT0	no sample	no sample	AMQSAPTC

Table 72. Sample programs demonstrating use of the MQI (C and COBOL) on IBM i (continued)

Technique	C (source) (1)	COBOL (source) (2)	RPG (source) (3)	Client (C executable)(4)
Using the publish/subscribe interface	AMQSPUBA, AMQSSUBA, AMQSSBXA	no sample	no sample	AMQSPUBC, AMQSSUBC, AMQSSBXC
Reconnectable clients (5)	AMQSPHAC, AMQSGHAC, AMQSMHAC	no sample	no sample	no sample
Using message consumers to asynchronously consume messages from multiple queues (5)	AMQSCBFO	no sample	no sample	no sample
Specifying TLS connection information on MQCONN	AMQSSSLC	no sample	no sample	AMQSSSLC
Connecting to the queue manager using MQCONN	AMQSCNXC	no sample	no sample	AMQSCNXC

Notes:

1. Source for the C samples is in the file QMQMSAMP/QCSRC. Include files exist as members in the file QMQM/H.
2. Source for the COBOL samples are in the files QMQMSAMP/QCBLLESRC. The members are named AMQ0 xxx 4, where xxx indicates the sample function.
3. Source for the RPG samples is in QMQMSAMP/QRPGLESRC. Members are named AMQ3 xxx 4, where xxx indicates the sample function. Copy members exist in QMQM/QRPGLESRC. Each member name has the suffix G.
4. The executable version of the IBM MQ MQI client samples share the same source as the samples that run in a server environment. Source for the samples in the client environment is same as the server. IBM MQ MQI client samples are linked with client library LIBMQIC and IBM MQ server samples are linked with server library LIBMQM.
5. If client executable for sample application of Reconnectable client and asynchronously consumer application has to be run, it has to be compiled and linked with threaded library LIBMQIC_R. Hence, it has to be run in threaded environment. Set the environment variable QIBM_MULTI_THREADED to 'Y' and run the application from qsh.

See Setting up IBM MQ with Java and JMS for more information.

In addition to these, the IBM MQ for IBM i sample option includes a sample data file, which you use as input to the sample programs, AMQSDATA and sample CL programs that demonstrate administration tasks. The CL samples are described in the Administering IBM i. You could use the sample CL program amqsamp4 to create queues to use with the sample programs described in this topic.

Preparing and running the sample programs:

Windows

UNIX

IBM i

After completing some initial preparation, you can then run the sample programs.

About this task

Before you run the sample programs, you must first create a queue manager and also create the queues that you need. You might also need to do some additional preparation, for example, if you want to run COBOL samples. After completing the necessary preparation, you can then run the sample programs.

Procedure

For information about how to prepare and run the sample programs, see the following topics:

- “Preparing and running sample programs on IBM i” on page 472
- “Preparing and running sample programs on UNIX and Linux” on page 473
- “Preparing and running sample programs on Windows” on page 474

Configuring a queue manager to accept client connections on Windows, UNIX and IBM i:

Windows

UNIX

IBM i

Before you can run the sample applications, you must first create a queue manager. You can then configure the queue manager to securely accept incoming connection requests from applications that are running in client mode.

Before you begin

Ensure the queue manager already exists and has been started. Determine whether channel authentication records are already enabled by issuing the MQSC command:

```
DISPLAY QMGR CHLAUTH
```

Important: This task expects that channel authentication records are enabled. If this is a queue manager used by other users and applications, changing this setting will affect all other users and applications. If your queue manager does not make use of channel authentication records then step 4 can be replaced with an alternate authentication method (for example, a security exit) which sets the MCAUSER to the *non-privileged-user-id* you will obtain in step 1.



You must know which channel name your application expects to use so that the application can be permitted to use the channel. You must also know which objects, for example queues or topics, your application expects to use so that your application can be permitted to use them.


About this task

This task creates a non-privileged user ID to be used for a client application which connects to the queue manager. Access is granted for the client application only to be able to use the channel it needs and the queue it needs by use of this user ID.

Procedure

1. Obtain a user ID on the system your queue manager is running on. For this task this user ID must not be a privileged administrative user. This user ID will be the authority under which the client connection will run on the queue manager.
2. Start a listener program with the following commands where:
 - qmgr-name* is the name of your queue manager
 - nnnn* is your chosen port number

- a.   For UNIX and Windows systems:

```
runmqtsr -t tcp -m qmgr-name -p nnnn
```
- b.  For IBM i:

```
STRMQLSR MQMNAME(qmgr-name) PORT(nnnn)
```
3. If your application uses the SYSTEM.DEF.SVRCONN then this channel is already defined. If your application uses another channel, create it by issuing the MQSC command:

```
DEFINE CHANNEL(' channel-name ') CHLTYPE(SVRCONN) TRPTYPE(TCP) +  
DESCR('Channel for use by sample programs')
```

channel-name is the name of your channel.
4. Create a channel authentication rule allowing only the IP address of your client system to use the channel by issuing the MQSC command:





```
SET CHLAUTH(' channel-name ') TYPE(ADDRESSMAP) ADDRESS(' client-machine-IP-address ') +  
MCAUSER(' non-privileged-user-id ')
```

channel-name is the name of your channel.
client-machine-IP-address is the IP address of your client system.





If your sample client application is running on the same machine as the queue manager then use an IP address of '127.0.0.1' if your application is going to connect using 'localhost'. If several different client machines are going to connect in, you can use a pattern or a range instead of a single IP address. See Generic IP addresses for details.

non-privileged-user-id is the user ID you obtained in step 1 on page 470
5. If your application uses the SYSTEM.DEFAULT.LOCAL.QUEUE then this queue is already defined. If your application uses another queue, create it by issuing the MQSC command:





```
DEFINE QLOCAL(' queue-name ') DESCR('Queue for use by sample programs')
```

queue-name is the name of your queue.
6. Grant access to connect to and inquire the queue manager:
 - a.    For  IBM i, UNIX and Windows systems issue the MQSC commands:

```
SET AUTHREC OBJTYPE(QMGR) PRINCIPAL(' non-privileged-user-id ') +  
AUTHADD(CONNECT, INQ)
```

non-privileged-user-id is the user ID you obtained in step 1 on page 470
7. If your application is a point-to-point application, that is it makes use of queues, grant access to allow inquiring and the putting and getting messages using your queue by the user ID to be used, by issuing the MQSC commands:
 - a.    For  IBM i, UNIX and Windows systems issue the MQSC commands:

```
SET AUTHREC PROFILE(' queue-name ') OBJTYPE(Queue) +  
PRINCIPAL(' non-privileged-user-id ') AUTHADD(PUT, GET, INQ, BROWSE)
```

queue-name is the name of your queue.
non-privileged-user-id is the user ID you obtained in step 1 on page 470
8. If your application is a publish/subscribe application, that is it makes use of topics, grant access to allow publishing and subscribing using your topic by the user ID to be used, by issuing the MQSC commands:
 - a.    For  IBM i, UNIX and Windows systems issue the MQSC commands:

```
SET AUTHREC PROFILE('SYSTEM.BASE.TOPIC') OBJTYPE(TOPIC) +  
PRINCIPAL(' non-privileged-user-id ') AUTHADD(PUB, SUB)
```

non-privileged-user-id is the user ID you obtained in step 1 on page 470

This will give *non-privileged-user-id* access to any topic in the topic tree, alternatively, you can define a topic object using **DEFINE TOPIC** and grant accesses only to the part of the topic tree referenced by that topic object. See Controlling user access to topics for details.

What to do next

Your client application can now connect to the queue manager and put or get messages using the queue.

Related information:

SET CHLAUTH


DEFINE CHANNEL

DEFINE QLOCAL

SET AUTHREC

 IBM MQ authorities on IBM i

Giving access to an IBM MQ object on UNIX or Linux systems and Windows

Preparing and running sample programs on IBM i: 

Before you run the sample programs on IBM i, you must first create a queue manager and also create the queues that you need. If you want to run COBOL samples, you might need to do some additional preparation.

About this task

The source for the IBM MQ for IBM i sample programs is provided in library QMQMSAMP as members of QCSRC, QCLSRC, QCBLESRC, and QRPGLSRC.

You can use your own queues when you run the samples, or you can run the sample program AMQSAMP4 to create some sample queues. The source for the AMQSAMP4 program is included in file QCLSRC in library QMQMSAMP. You can compile it by using the CRTCLPGM command.

To run the samples, either use the C executable versions, which are supplied in the library QMQM, or compile them in a similar way to any other IBM MQ application.

Procedure

1. Create a queue manager and set up the default definitions. You must do this before you can run any of the sample programs. For more information about creating a queue manager, see *Administering IBM MQ*. For information about configuring a queue manager to securely accept incoming connection requests from applications that are running in client mode, see “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470.

2. To call one of the sample programs by using data from member PUT in file AMQSDATA of library QMQMSAMP, use a command like:

```
CALL PGM(QMQM/AMQSPUT4) PARM('QMQMSAMP/AMQSDATA(PUT)')
```

Note: For a compiled module to use the IFS file system, specify the option SYSIFCOPT(*IFSIO) on CRTCMOD, then the file name, passed as a parameter, must be specified in the following format:
home/me/myfile

3. If you want to use the COBOL versions of the Inquire, Set, and Echo examples, change the process definitions before you run these samples. For the Inquire, Set, and Echo examples, the sample definitions trigger the C versions of these samples. If you want the COBOL versions, you must change the process definitions:
 - SYSTEM.SAMPLE.INQPROCESS
 - SYSTEM.SAMPLE.SETPROCESS

- SYSTEM.SAMPLE.ECHOPROCESS

On IBM i, you can use the **CHGMQMPCR** command (for details, see Change MQ Process (CHGMQMPCR)), or edit and run the **AMQSAMP4** command with the alternative definition.


4. Run the sample programs. For more information on the parameters that each of the samples expects, see the descriptions of the individual samples.

Note: For the COBOL sample programs, when you pass queue names as parameters, you must provide 48 characters, padding with blank characters if necessary. Anything other than 48 characters causes the program to fail with reason code 2085.

Related reference:

“Samples for IBM i” on page 468

The techniques demonstrated by the sample programs for IBM MQ on IBM i systems.

Preparing and running sample programs on UNIX and Linux: 

Before you run the sample programs on UNIX, you must first create a queue manager and also create the queues that you need. If you want to run COBOL samples, you might need to do some additional preparation.

About this task

The IBM MQ on UNIX and Linux systems sample files are in the directories listed in Table 73 if the defaults were used at installation time.

Table 73. Where to find the samples for IBM MQ on UNIX and Linux systems

Content	Directory
source files	<i>MQ_INSTALLATION_PATH</i> /samp
dead-letter queue handler source files	<i>MQ_INSTALLATION_PATH</i> /samp/dlq
executable files	<i>MQ_INSTALLATION_PATH</i> /samp/bin

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

The samples need a set of queues to work with. You can either use your own queues or run the sample MQSC file *amqscos0.tst* to create a set. To run the samples, either use the executable versions supplied or compile the source versions as you would any other applications, using an ANSI compiler.

Procedure

1. Create a queue manager and set up the default definitions. You must do this before you can run any of the sample programs. For more information about creating a queue manager, see Administering IBM MQ. For information about configuring a queue manager to securely accept incoming connection requests from applications that are running in client mode, see “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470.
2. If you are not using your own queues, run the sample MQSC file *amqscos0.tst* to create a set of queues. To do this on UNIX and Linux systems, enter:

```
runmqsc QManagerName <amqscos0.tst > /tmp/sampobj.out
```

Check the *sampobj.out* file to ensure that there are no errors.

3. If you want to use the COBOL versions of the Inquire, Set, and Echo examples, change the process definitions before you run these samples. For the Inquire, Set, and Echo examples, the sample definitions trigger the C versions of these samples. If you want the COBOL versions, you must change the process definitions:

- SYSTEM.SAMPLE.INQPROCESS

- SYSTEM.SAMPLE.SETPROCESS
- SYSTEM.SAMPLE.ECHOPROCESS

On Windows, do this by editing the `amqscos0.tst` file and changing the C executable file names to the COBOL executable file names before using the `runmqsc` command to run these samples.

4. Run the sample programs. To run a sample, enter its name followed by any parameters, for example:

```
amqsput myqueue qmanagername
```

where *myqueue* is the name of the queue on which the messages are going to be put, and *qmanagername* is the queue manager that owns *myqueue*.

For more information on the parameters that each of the samples expects, see the descriptions of the individual samples.

Note: For the COBOL sample programs, when you pass queue names as parameters, you must provide 48 characters, padding with blank characters if necessary. Anything other than 48 characters causes the program to fail with reason code 2085.

Related reference:

“Samples for UNIX and Linux systems” on page 459

The techniques demonstrated by the sample programs for IBM MQ on UNIX and Linux systems.

Preparing and running sample programs on Windows: **Windows**

Before you run the sample programs on Windows, you must first create a queue manager and also create the queues that you need. If you want to run COBOL samples, you might need to do some additional preparation.

About this task

The IBM MQ for Windows sample files are in the directories listed in Table 74, if the defaults were used at installation time. The installation drive defaults to <c:>.

Table 74. Where to find the samples for IBM MQ for Windows

Content	Directory
C source code	<code>MQ_INSTALLATION_PATH\Tools\C\Samples</code>
Source code for dead-letter handler sample	<code>MQ_INSTALLATION_PATH\Tools\C\Samples\DLQ</code>
COBOL source code	<code>MQ_INSTALLATION_PATH\Tools\Cobol\Samples</code>
C executable files ¹	<code>MQ_INSTALLATION_PATH\Tools\C\Samples\Bin (32-bit versions)</code> <code>MQ_INSTALLATION_PATH\Tools\C\Samples\Bin64 (64-bit versions)</code>
Sample MQSC files	<code>MQ_INSTALLATION_PATH\Tools\MQSC\Samples</code>
Visual Basic source code	<code>MQ_INSTALLATION_PATH\Tools\VB\SampVB6</code>
.NET samples	<code>MQ_INSTALLATION_PATH\Tools\dotnet\Samples</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Note: 64-bit versions are available of some C executable file samples.

The samples need a set of queues to work with. You can either use your own queues or run the sample MQSC file `amqscos0.tst` to create a set of queues. To run the samples, either use the executable versions supplied or compile the source versions as you would any other IBM MQ for Windows applications.

Procedure

1. Create a queue manager and set up the default definitions. You must do this before you can run any of the sample programs. For more information about creating a queue manager, see *Administering IBM MQ*. For information about configuring a queue manager to securely accept incoming connection requests from applications that are running in client mode, see “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470.
2. If you are not using your own queues, run the sample MQSC file `amqscos0.tst` to create a set of queues. To do this on Windows systems enter:

```
runmqsc QManagerName < amqscos0.tst > sampobj.out
```

Check the `sampobj.out` file to ensure that there are no errors. This file is in your current directory.

Note:

3. If you want to use the COBOL versions of the Inquire, Set, and Echo examples, change the process definitions before you run these samples. For the Inquire, Set, and Echo examples, the sample definitions trigger the C versions of these samples. If you want the COBOL versions, you must change the process definitions:

- SYSTEM.SAMPLE.INQPROCESS
- SYSTEM.SAMPLE.SETPROCESS
- SYSTEM.SAMPLE.ECHOPROCESS

On Windows, do this by editing the `amqscos0.tst` file and changing the C executable file names to the COBOL executable file names before using the `runmqsc` command to run these samples.

4. Run the sample programs. To run a sample, enter its name followed by any parameters, for example:

```
amqsput myqueue qmanagername
```

where *myqueue* is the name of the queue on which the messages are going to be put, and *qmanagername* is the queue manager that owns *myqueue*.

For more information on the parameters that each of the samples expects, see the descriptions of the individual samples.

Note: For the COBOL sample programs, when you pass queue names as parameters, you must provide 48 characters, padding with blank characters if necessary. Anything other than 48 characters causes the program to fail with reason code 2085.

Related reference:

“Samples for IBM MQ for Windows” on page 465

The techniques demonstrated by the sample programs for IBM MQ for Windows.

“Visual Basic samples for IBM MQ for Windows” on page 467

The techniques demonstrated by the sample programs for IBM MQ on Windows systems.

The API exit sample program:

The sample API exit generates an MQI trace to a user-specified file with a prefix defined in the `MQAPI_TRACE_LOGFILE` environment variable.

For more information about API exits, see “Writing and compiling API exits” on page 332.

Source

```
amqsaxe0.c
```

Binary

```
amqsaxe
```

Configuring for the sample exit

1. Add the following to the `qm.ini` file.

Platforms other than Windows

```
ApiExitLocal:  
Sequence=100  
Function=EntryPoint  
Module= MQ_INSTALLATION_PATH/samp/bin/amqsaxe  
Name=SampleApiExit
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM MQ is installed.

Windows

```
ApiExitLocal:  
Sequence=100  
Function=EntryPoint  
Module= MQ_INSTALLATION_PATH\Tools\c\Samples\bin\amqsaxe  
Name=SampleApiExit
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM MQ is installed.

2. Set the environment variable

```
MQAPI_TRACE_LOGFILE=/tmp/MqiTrace
```

3. Run your application.

Output files are created in the /tmp directory with names like: *MqiTrace.pid.tid.log*

The Asynchronous consumption sample program:

The *amqscbf* sample program demonstrates the use of MQCB and MQCTL to consume messages from multiple queues asynchronously.

amqscbf is provided as C source code, and a binary client and server executable on Windows, UNIX and Linux platforms.

The program is started from the command line and takes the following optional parameters:

Usage: [Options] **Queue Name** {*queue_name*}

where Options are:

```
-m Queue Manager Name  
-o Open options  
-r Reconnect Type  
  d Reconnect Disabled  
  r Reconnect  
  m Reconnect Queue Manager
```

Provide more than one queue name to read messages from multiple queues (a maximum of ten queues are supported by the sample.)

Note: **Reconnect type** is only valid for client programs.

Example

The example shows *amqscbf* run as a server program reading one message from QL1 and then being stopped.

Use IBM MQ Explorer to put a test message on QL1. Stop the program by pressing enter.

```
C:\>amqscbf QL1  
Sample AMQSCBF0 start
```

```
Press enter to end  
Message Call (9 Bytes) :  
Message 1
```

```
Sample AMQSCBF0 end
```

What amqscbf demonstrates

The sample shows how to read messages from multiple queues in the order of their arrival. This would require a lot more code using synchronous MQGET. In the case of asynchronous consumption, no polling is required, and thread and storage management is performed by IBM MQ. A "real world" example would need to deal with errors; in the sample errors are written out to the console.

The sample code has the following steps,

1. Define the single message consumption callback function,

```
void MessageConsumer(MQHCONN      hConn,
                    MQMD          * pMsgDesc,
                    MQGMO         * pGetMsgOpts,
                    MQBYTE        * Buffer,
                    MQCBC         * pContext)
{ ... }
```

2. Connect to the queue manager,

```
MQCONN(XQMName, &cno, &Hcon, &CompCode, &CReason);
```

3. Open the input queues, and associate each one with the MessageConsumer callback function,

```
MQOPEN(Hcon, &od, 0_options, &Hobj, &OpenCode, &Reason);
cbd.CallbackFunction = MessageConsumer;
MQCB(Hcon, MQOP_REGISTER, &cbd, Hobj, &md, &gmo, &CompCode, &Reason);
```

cbd.CallbackFunction does not need to be set for each queue; it is an input-only field. But you could associate a different callback function with each queue.

4. Start consumption of the messages,

```
MQCTL(Hcon, MQOP_START, &ctlo, &CompCode, &Reason);
```

5. Wait until the user has pressed enter and then stop consumption of messages,

```
MQCTL(Hcon, MQOP_STOP, &ctlo, &CompCode, &Reason);
```

6. Finally disconnect from the queue manager,

```
MQDISC(&Hcon, &CompCode, &Reason);
```

The Asynchronous Put sample program:

Learn about running the amqsapt sample and the design of the Asynchronous Put sample program.

The asynchronous put sample program puts messages on a queue using the asynchronous MQPUT call and then retrieves status information using the MQSTAT call. See "Features demonstrated in the sample programs on Multiplatforms" on page 459 for the name of this program on different platforms.

Running the amqsapt sample

This program takes up to 6 parameters:

1. The name of the target queue (required)
2. The name of the queue manager (optional)
3. Open options (optional)
4. Close options (optional)
5. The name of the target queue manager (optional)
6. The name of the dynamic queue (optional)

If a queue manager is not specified, amqsapt connects to the default queue manager.

Design of the Asynchronous Put sample program

The program uses the MQOPEN call with the output options supplied, or with the MQOO_OUTPUT and MQOO_FAIL_IF_QUIESCING options to open the target queue for putting messages.

If it cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call. To keep the program simple, on this and on subsequent MQI calls, the program uses default values for many of the options.

For each line of input, the program reads the text into a buffer and uses the MQPUT call with MQPMO_ASYNC_RESPONSE to create a datagram message containing the text of that line and asynchronously put it to the target queue. The program continues until it reaches the end of the input or the MQPUT call fails. If the program reaches the end of the input, it closes the queue using the MQCLOSE call.

The program then issues the MQSTAT call, returning an MQSTS structure, and displays messages containing the number of messages put successfully, the number of messages put with a warning, and the number of failures.

The Browse sample programs:

The Browse sample programs browse messages on a queue using the MQGET call.

See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

Design of the Browse sample program

The program opens the target queue using the MQOPEN call with the MQOO_BROWSE option. If it cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call.

For each message on the queue, the program uses the MQGET call to copy the message from the queue, then displays the data contained in the message. The MQGET call uses these options:

MQGMO_BROWSE_NEXT

After the MQOPEN call, the browse cursor is positioned logically before the first message in the queue, so this option causes the *first* message to be returned when the call is first made.

MQGMO_NO_WAIT

The program does not wait if there are no messages on the queue.

MQGMO_ACCEPT_TRUNCATED_MSG

The MQGET call specifies a buffer of fixed size. If a message is longer than this buffer, the program displays the truncated message, together with a warning that the message has been truncated.

The program demonstrates how you must clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call, because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The program continues to the end of the queue; the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code and the program displays a warning message. If the MQGET call fails, the program displays an error message that contains the reason code.

The program then closes the queue using the MQCLOSE call.

The Browse sample programs for UNIX, Linux, and Windows:

Consider using this topic when learning about Browse sample programs on UNIX, Linux, and Windows.

The C version of the program takes 2 parameters

1. The name of the source queue (necessary)
2. The name of the queue manager (optional)

If a queue manager is not specified, it connects to the default one. For example, enter one of the following:

- `amqsgbr myqueue qmanagername`
- `amqsgbr0 myqueue`

where `myqueue` is the name of the queue that the messages will be viewed from, and `qmanagername` is the queue manager that owns `myqueue`.

If you omit the `qmanagername`, when running the C sample, it assumes that the default queue manager owns the queue.

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

Please enter the name of the target queue

Only the first 50 characters of each message are displayed, followed by - - - truncated when this is the case.

The Browse sample programs on IBM i:

Each program retrieves copies of all the messages on the queue that you specify when you call the program; the messages remain on the queue.

You can use the supplied queue `SYSTEM.SAMPLE.LOCAL`; run the Put sample program first to put some messages on the queue. You can use the queue `SYSTEM.SAMPLE.ALIAS`, which is an alias name for the same local queue. The program continues until it reaches the end of the queue or an MQI call fails.

The C samples let you specify the queue manager name, generally as the second parameter, in a similar fashion to the Windows systems samples. For example:

```
CALL PGM(QMQM/AMQSTRG4) PARM('SYSTEM.SAMPLE.TRIGGER' 'QM01')
```

If a queue manager is not specified, it connects to the default one. This is also relevant to the RPG samples. However, with the RPG samples you must supply a queue manager name rather than allowing it to default.

The Browser sample program:

The Browser sample program reads and writes both the message descriptor and the message content fields of all the messages on a queue.

The sample program is written as a utility, not just to demonstrate a technique. See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

This program takes these positional parameters:

1. The name of the source queue (required)
2. The name of the queue manager (required)
3. An optional parameter for properties (optional)

These programs also use an environment variable named **MQSAMP_USER_ID** which should be set to the user ID to be used for connection authentication. When this is set, the program prompts for a password to accompany that user ID.

To run these programs, enter one of the following commands:

- `amqsbcg myqueue qmanagername`
- `amqsbcg myqueue qmanagername`

where *myqueue* is the name of the queue on which the messages are going to be browsed, and *qmanagername* is the queue manager that owns *myqueue*.

It reads each message from the queue and writes the following to stdout:

- Formatted message descriptor fields
- Message data (dumped in hex and, where possible, character format)

Table 75. Permissible values for the property parameter

Value	Behavior
0	Default behavior, as it was for V6. The properties that get delivered to the application depend on the PropertyControl queue attribute that the message is retrieved from.
1	A message handle is created and used with the MQGET. Properties of the message, except those contained in the message descriptor (or extension) are displayed in a similar fashion to the message descriptor. For example: <pre>****Message properties**** property name: property value</pre> Or if no properties are available: <pre>****Message properties**** None</pre> Numeric values are displayed using <code>printf</code> , string values are surrounded in single quotation marks, and byte strings are surrounded with X and single quotation marks, as for the message descriptor.
2	MQGMO_NO_PROPERTIES is specified, so that only message descriptor properties will be returned.
3	MQGMO_PROPERTIES_FORCE_MQRFH2 is specified, so that all properties are returned in the message data.
4	MQGMO_PROPERTIES_COMPATIBILITY is specified, so that all properties can be returned depending on whether a Version 6 property is included, otherwise the properties are discarded.

The program is restricted to printing the first 65535 characters of the message, and fails with the reason *truncated msg* if a longer message is read.

For an example of the output from this utility, see .

The CICS transaction sample:

A sample CICS transaction program is provided, named `amqscic0.ccs` for source code and `amqscic0` for the executable version. You can build transactions using the standard CICS facilities.

See “Building a procedural application” on page 382 for details on the commands needed for your platform.

The transaction reads messages from the transmission queue `SYSTEM.SAMPLE.CICS.WORKQUEUE` on the default queue manager and places them onto the local queue, the name of which is contained in the transmission header of the message. Any failures are sent to the queue `SYSTEM.SAMPLE.CICS.DLQ`.

Note: You can use a sample MQSC script `amqscic0.tst` to create these queues and sample input queues.

The Connect sample program:

The Connect sample program allows you to explore the `MQCONN` call and its options from a client. The sample connects to the queue manager using the `MQCONN` call, inquires about the name of the queue manager using the `MQINQ` call, and displays it. Also, learn about running the `amqscnxc` sample.

Note: The Connect sample program is a client sample. You can compile and run it on a server but the function is meaningful only on a client, and only client-executable files are supplied.

Running the `amqscnxc` sample

The command-line syntax of the Connect sample program is:

```
amqscnxc [-x ConnName [-c SvrconnChannelName]] [-u User] [QMgrName]
```

The parameters are optional and their order is not important except for `QMgrName`, which, if specified, must come last. The parameters are:

ConnName

The TCP/IP connection name of the server queue manager

If you do not specify the TCP/IP connection name, `MQCONN` is issued with the `ClientConnPtr` set to `NULL`.

SvrconnChannelName

The name of the server connection channel

If you specify the TCP/IP connection name but not the server connection channel (the reverse is not allowed), the sample uses the name `SYSTEM.DEF.SVRCONN`.

User The user name to be used for connection authentication

If you specify this the program will prompt for a password to accompany that user ID.

QMgrName

The name of the target queue manager

If you do not specify the target queue manager, the sample connects to whichever queue manager is listening at the given TCP/IP connection name.

Note: If you enter a question mark as the only parameter, or if you enter incorrect parameters, you get a message explaining how to use the program.

If you run the sample with no command-line options, the contents of the MQSERVER environment variable are used to determine the connection information. (In this example MQSERVER is set to SYSTEM.DEF.SVRCONN/TCP/machine.site.company.com.) You see output like this:

```
Sample AMQSCNXC start
Connecting to the default queue manager
with no client connection information specified.
Connection established to queue manager machine
```

Sample AMQSCNXC end

If you run the sample and provide a TCP/IP connection name and a server connection channel name but no target queue manager name, like this:

```
amqscnxc -x machine.site.company.com -c SYSTEM.ADMIN.SVRCONN
```

the default queue manager name is used and you see output like this:

```
Sample AMQSCNXC start
Connecting to the default queue manager
using the server connection channel SYSTEM.ADMIN.SVRCONN
on connection name machine.site.company.com.
Connection established to queue manager MACHINE
```

Sample AMQSCNXC end

If you run the sample and provide a TCP/IP connection name and a target queue manager name, like this:

```
amqscnxc -x machine.site.company.com MACHINE
```

you see output like this:

```
Sample AMQSCNXC start
Connecting to queue manager MACHINE
using the server connection channel SYSTEM.DEF.SVRCONN
on connection name machine.site.company.com.
Connection established to queue manager MACHINE
```

Sample AMQSCNXC end

The Data-Conversion sample program:

The data-conversion sample program is a skeleton of a data conversion exit routine. Learn about the design of the data-conversion sample.

See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

Design of the data-conversion sample

Each data-conversion exit routine converts a single named message format. This skeleton is intended as a wrapper for code fragments generated by the data-conversion exit generation utility program.

The utility produces one code fragment for each data structure; several such structures make up a format, so several code fragments are added to this skeleton to produce a routine to do data conversion of the entire format.

The program then checks whether the conversion is a success or failure, and returns the values required to the caller.

Database coordination samples:

Two samples are provided that demonstrate how IBM MQ can coordinate both IBM MQ updates and database updates within the same unit of work.

These samples are:

1. AMQSXAS0 (in C) or AMQ0XAS0 (in COBOL), which updates a single database within an IBM MQ unit of work.
2. AMQSXAG0 (in C) or AMQ0XAG0 (in COBOL), AMQSXAB0 (in C) or AMQ0XAB0 (in COBOL), and AMQSXAF0 (in C) or AMQ0XAF0 (in COBOL), which together update two databases within an IBM MQ unit of work, showing how multiple databases can be accessed. These samples are provided to show the use of the MQBEGIN call, mixed SQL and IBM MQ calls, and where and when to connect to a database.

Figure 83 shows how the samples provided are used to update databases:

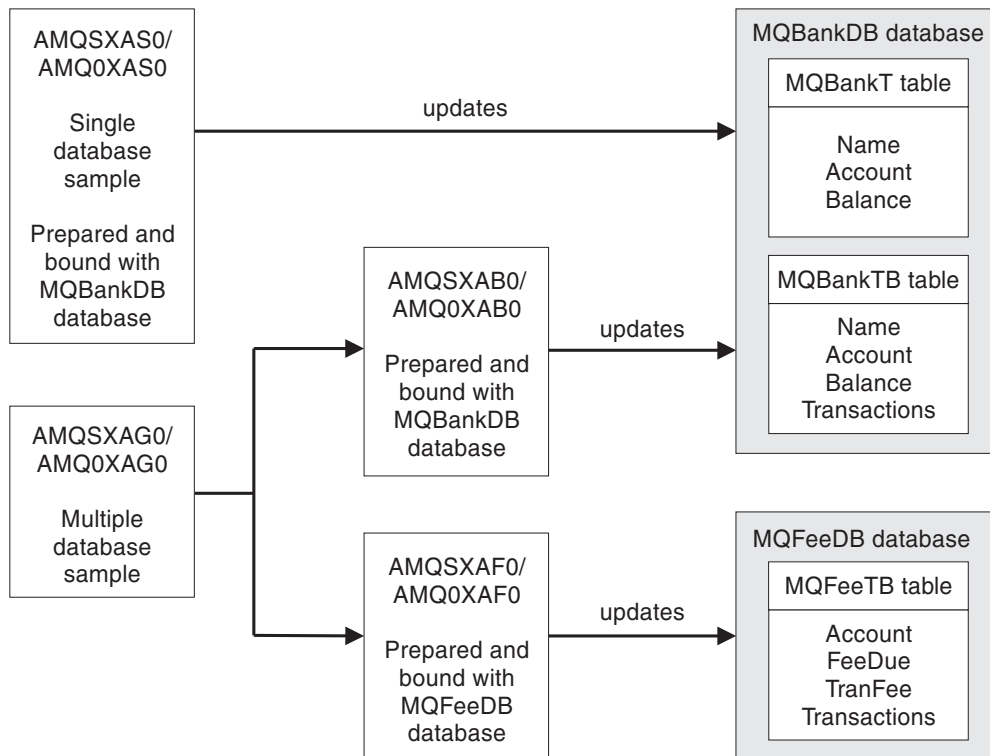


Figure 83. The database coordination samples

The programs read a message from a queue (under syncpoint), then, using the information in the message, obtain the relevant information from the database and update it. The new status of the database is then printed.

The program logic is as follows:

1. Use the name of the input queue from the program argument
2. Connect to the default queue manager (or optionally the supplied name in C) using MQCONN
3. Open a queue (using MQOPEN) for input while there are no failures
4. Start a unit of work using MQBEGIN
5. Get the next message (using MQGET) from the queue under syncpoint

6. Get information from databases
7. Update information from databases
8. Commit changes using MQCMIT
9. Print updated information (no message being available counts as a failure, and the loop ends)
10. Close the queue using MQCLOSE
11. Disconnect from the queue using MQDISC

SQL cursors are used in the samples, so that reads from the databases (that is, multiple instances) are locked while a message is being processed, allowing multiple instances of these programs to run simultaneously. The cursors are explicitly opened, but implicitly closed by the MQCMIT call.

The single database sample (AMQXSAS0 or AMQ0XAS0) has no SQL CONNECT statements and the connection to the database is implicitly made by IBM MQ with the MQBEGIN call. The multiple database sample (AMQXSAG0 or AMQ0XAG0, AMQSXAB0 or AMQ0XAB0, and AMQSXAF0 or AMQ0XAF0) has SQL CONNECT statements, as some database products allow only one active connection. If this is not the case for your database product, or if you are accessing a single database in multiple database products, the SQL CONNECT statements can be removed.

The samples are prepared with the IBM Db2 database product, so you might need to modify them to work with other database products.

The SQL error checking uses routines in UTIL.C and CHECKERR.CBL supplied by Db2. These must be compiled or replaced before compiling and linking.

Note: If you are using the Micro Focus COBOL source CHECKERR.MFC for SQL error checking, you must change the program ID to uppercase, that is CHECKERR, for AMQ0XAS0 to link correctly.

Creating the databases and tables:

Create the databases and tables before compiling the samples.

To create the databases, use the usual method for your database product, for example:

```
DB2 CREATE DB MQBankDB
DB2 CREATE DB MQFeeDB
```

Create the tables using SQL statements as follows:

In C:

```
EXEC SQL CREATE TABLE MQBankT(Name          VARCHAR(40) NOT NULL,
                               Account       INTEGER    NOT NULL,
                               Balance       INTEGER    NOT NULL,
                               PRIMARY KEY (Account));

EXEC SQL CREATE TABLE MQBankTB(Name         VARCHAR(40) NOT NULL,
                                Account      INTEGER    NOT NULL,
                                Balance      INTEGER    NOT NULL,
                                Transactions  INTEGER,
                                PRIMARY KEY (Account));

EXEC SQL CREATE TABLE MQFeeTB(Account      INTEGER    NOT NULL,
                                FeeDue      INTEGER    NOT NULL,
                                TranFee     INTEGER    NOT NULL,
                                Transactions INTEGER,
                                PRIMARY KEY (Account));
```

In COBOL:

```
EXEC SQL CREATE TABLE
MQBankT(Name          VARCHAR(40) NOT NULL,
        Account       INTEGER    NOT NULL,
        Balance       INTEGER    NOT NULL,
        PRIMARY KEY (Account))
END-EXEC.
```

```
EXEC SQL CREATE TABLE
MQBankTB(Name          VARCHAR(40) NOT NULL,
        Account       INTEGER    NOT NULL,
        Balance       INTEGER    NOT NULL,
        Transactions  INTEGER,
        PRIMARY KEY (Account))
END-EXEC.
```

```
EXEC SQL CREATE TABLE
MQFeeTB(Account       INTEGER    NOT NULL,
        FeeDue        INTEGER    NOT NULL,
        TranFee       INTEGER    NOT NULL,
        Transactions  INTEGER,
        PRIMARY KEY (Account))
END-EXEC.
```

Enter data into the tables using SQL statements as follows:

```
EXEC SQL INSERT INTO MQBankT VALUES ('Mr Fred Bloggs',1,0);
EXEC SQL INSERT INTO MQBankT VALUES ('Mrs S Smith',2,0);
EXEC SQL INSERT INTO MQBankT VALUES ('Ms Mary Brown',3,0);
:
:
EXEC SQL INSERT INTO MQBankTB VALUES ('Mr Fred Bloggs',1,0,0);
EXEC SQL INSERT INTO MQBankTB VALUES ('Mrs S Smith',2,0,0);
EXEC SQL INSERT INTO MQBankTB VALUES ('Ms Mary Brown',3,0,0);
:
:
EXEC SQL INSERT INTO MQFeeTB VALUES (1,0,50,0);
EXEC SQL INSERT INTO MQFeeTB VALUES (2,0,50,0);
EXEC SQL INSERT INTO MQFeeTB VALUES (3,0,50,0);
:
:
```

Note: For COBOL, use the same SQL statements but add END_EXEC at the end of each line.

Precompiling, compiling, and linking the samples:

Learn about precompiling, compiling, and linking samples in C and COBOL.

Precompile the .SQC files (in C) and .SQB files (in COBOL), and bind them against the appropriate database to produce the .C or .CBL files. To do this, use the typical method for your database product.

Precompiling in C

```
db2 connect to MQBankDB
db2 prep AMQSXAS0.SQC
db2 connect reset
```

```
db2 connect to MQBankDB
db2 prep AMQSXAB0.SQC
db2 connect reset
```

```
db2 connect to MQFeeDB
db2 prep AMQSXAF0.SQC
db2 connect reset
```

Precompiling in COBOL

```
db2 connect to MQBankDB
db2 prep AMQ0XAS0.SQB bindfile target ibmcob
db2 bind AMQ0XAS0.BND
```

```
db2 connect reset
```

```
db2 connect to MQBankDB
db2 prep AMQ0XAB0.SQB bindfile target ibmcob
db2 bind AMQ0XAB0.BND
db2 connect reset
```

```
db2 connect to MQFeeDB
db2 prep AMQ0XAF0.SQB bindfile target ibmcob
db2 bind AMQ0XAF0.BND
db2 connect reset
```

Compiling and linking

The following sample commands use the symbols *DB2TOP* and *MQ_INSTALLATION_PATH*. *DB2TOP* represents the installation directory for the Db2 product. *MQ_INSTALLATION_PATH* represents the high-level directory in which IBM MQ is installed.

- On AIX, the directory path is:
`/usr/lpp/db2_05_00`
- On HP-UX and Solaris, the directory path is:
`/opt/IBMd2/V5.0`
- On Windows systems, the directory path depends on the path chosen when installing the product. If you chose the default settings the path is:
`c:\sqllib`

Note: Before issuing the link command on Windows systems, ensure that the LIB environment variable contains paths to the Db2 and IBM MQ libraries.

Copy the following files into a temporary directory:

- The `amqsxag0.c` file from your IBM MQ installation

Note: This file can be found in the following directories:

- On UNIX and Linux systems:
`MQ_INSTALLATION_PATH/samp/xatm`
- On Windows systems:
`MQ_INSTALLATION_PATH\tools\c\samples\xatm`
- The `.c` files that you have obtained by precompiling the `.sqc` source files, `amqsxas0.sqc`, `amqsxaf0.sqc`, and `amqsxab0.sqc`
- The files `util.c` and `util.h` from your Db2 installation.

Note: These files can be found in the directory:

`DB2TOP/samples/c`

Build the object files for each `.c` file using the following compiler command for the platform that you are using:

- AIX
`xlc_r -I MQ_INSTALLATION_PATH/inc -I DB2TOP/include -c -o FILENAME.o FILENAME.c`
- HP-UX
`cc -Aa +z -I MQ_INSTALLATION_PATH/inc -I DB2TOP/include -c -o FILENAME.o FILENAME.c`
- Solaris

```
cc -Aa -KPIC -mt -I MQ_INSTALLATION_PATH
/inc -I DB2TOP/include -c -o
FILENAME.o FILENAME.c
```

- Windows systems

```
cl /c /I MQ_INSTALLATION_PATH\tools\c\include /I DB2TOP\include
FILENAME.c
```

Build the amqsxag0 executable file using the following link command for the platform that you are using:

- AIX

```
xlc_r -H512 -T512 -L DB2TOP/lib -ldb2 -L MQ_INSTALLATION_PATH/lib
-lmqm util.o amqsxaf0.o amqsxab0.o amqsxag0.o -o amqsxag0
```

- HP-UX Revision 11i

```
ld -E -L DB2TOP/lib -ldb2 -L MQ_INSTALLATION_PATH/lib -lmqm -lc -lpthread -lcl
/lib/crt0.o util.o amqsxaf0.o amqsxab0.o amqsxag0.o -o amqsxag0
```

- Solaris

```
cc -mt -L DB2TOP/lib -ldb2 -L MQ_INSTALLATION_PATH/lib
-lmqm -lthread -lsocket -lc -lnsl -ldl util.o
amqsxaf0.o amqsxab0.o amqsxag0.o -o amqsxag0
```

- Windows systems

```
link util.obj amqsxaf0.obj amqsxab0.obj amqsxag0.obj mqm.lib db2api.lib
/out:amqsxag0.exe
```

Build the amqsxas0 executable file using the following compiling and linking commands for the platform that you are using:

- AIX

```
xlc_r -H512 -T512 -L DB2TOP/lib -ldb2
-L MQ_INSTALLATION_PATH/lib -lmqm util.o amqsxas0.o -o amqsxas0
```

- HP-UX Revision 11i

```
ld -E -L DB2TOP/lib -ldb2 -L MQ_INSTALLATION_PATH/lib -lmqm -lc -lpthread
-lcl /lib/crt0.o util.o amqsxas0.o -o amqsxas0
```

- Solaris

```
cc -mt -L DB2TOP/lib -ldb2 -L MQ_INSTALLATION_PATH/lib
-lqm -lthread -lsocket -lc -lnsl -ldl util.o
amqsxas0.o -o amqsxas0
```

- Windows systems

```
link util.obj amqsxas0.obj mqm.lib db2api.lib /out:amqsxas0.exe
```

Additional information

If you are working on AIX or HP-UX and want to access Oracle, use the xlc_r compiler and link to libmqm_r.a.

Running the samples:

Use this information to learn how to configure the queue manager before running database coordination samples on C and COBOL.

Before you run the samples, configure the queue manager with the database product that you are using. For information about how to do this, see Scenario 1: Queue manager performs the coordination.

The following titles provide information about how to run samples in C and COBOL:

- "C samples"
- "COBOL samples"

C samples

Messages must be in the following format to be read from a queue:

```
UPDATE Balance change=nnn WHERE Account=nnn
```

AMQSPUT can be used to put the messages on the queue.

The database coordination samples take two parameters:

1. Queue name (required)
2. Queue manager name (optional)

Assuming that you have created and configured a queue manager for the single database sample called singDBQM, with a queue called singDBQ, you increment Mr Fred Bloggs's account by 50 as follows:

```
AMQSPUT singDBQ singDBQM
```

Then key in the following message:

```
UPDATE Balance change=50 WHERE Account=1
```

You can put multiple messages on the queue.

```
AMQSXAS0 singDBQ singDBQM
```

The updated status of Mr Fred Bloggs's account is then printed.

Assuming that you have created and configured a queue manager for the multiple-database sample called multDBQM, with a queue called multDBQ, you decrement Ms Mary Brown's account by 75 as follows:

```
AMQSPUT multDBQ multDBQM
```

Then key in the following message:

```
UPDATE Balance change=-75 WHERE Account=3
```

You can put multiple messages on the queue.

```
AMQSXAG0 multDBQ multDBQM
```

The updated status of Ms Mary Brown's account is then printed.

COBOL samples

Messages must be in the following format to be read from a queue:

```
UPDATE Balance change=snnnnnnnn WHERE Account=nnnnnnnn
```


For simplicity, the Balance change must be a signed eight-character number and the Account must be an eight-character number.

The sample AMQSPUT can be used to put the messages on the queue.

The samples take no parameters and use the default queue manager. It can be configured to run only one of the samples at any time. Assuming that you have configured the default queue manager for the single database sample, with a queue called singDBQ, you increment Mr Fred Bloggs's account by 50 as follows:

```
AMQSPUT singDBQ
```

Then key in the following message:

```
UPDATE Balance change=+00000050 WHERE Account=00000001
```

You can put multiple messages on the queue:

```
AMQ0XAS0
```

Type in the name of the queue:

```
singDBQ
```

The updated status of Mr Fred Bloggs's account is then printed.

Assuming that you have configured the default queue manager for the multiple database sample, with a queue called multDBQ, you decrement Ms Mary Brown's account by 75 as follows:

```
AMQSPUT multDBQ
```

Then key in the following message:

```
UPDATE Balance change=-00000075 WHERE Account=00000003
```

You can put multiple messages on the queue:

```
AMQ0XAG0
```

Type in the name of the queue:

```
multDBQ
```

The updated status of Ms Mary Brown's account is then printed.

Dead-letter queue handler sample:

A sample dead-letter queue handler is provided, the name of the executable version is amqsdq. If you want a dead-letter queue handler that is different from RUNMQDLQ, the source of the sample is available for you to use as your base.

The sample is similar to the dead-letter handler provided within the product but trace and error reporting are different. There are two environment variables available to you:

ODQ_TRACE

Set to YES or yes to switch tracing on

ODQ_MSG

Set to the name of the file containing error and information messages. The file provided is called amqsdq.msg.

You need to make these variables known to your environment using either the **export** or **set** commands, depending on your platform; trace is turned off using the **unset** command.

You can modify the error message file, `amqsdlq.msg`, to suit your own requirements. The sample puts messages to `stdout`, **not** to the IBM MQ error log file.

The Administering or the *System Management Guide* for your platform explains how the dead-letter handler works, and how you run it.

The Distribution List sample program:

The Distribution List sample `amqsptl0` gives an example of putting a message on several message queues. It is based on the MQPUT sample, `amqsput0`.

Running the Distribution List sample, `amqsptl0`

The Distribution List sample runs in a similar way to the Put samples.

It takes the following parameters:

- The names of the queues
- The names of the queue managers

These values are entered as pairs. For example:

```
amqsptl0 queue1 qmanagename1 queue2 qmanagename2
```

The queues are opened using MQOPEN and messages are put to the queues using MQPUT. Reason codes are returned if any of the queue or queue manager names are not recognized.

Remember to define channels between queue managers so that messages can flow between them. The sample program does not do that for you.

Design of the Distribution List sample

Put Message Records (MQPMRs) specify message attributes for each destination. The sample provides values for *MsgId* and *CorrelId*, and these override the values specified in the MQMD structure.

The *PutMsgRecFields* field in the MQPMO structure indicates which fields are present in the MQPMRs:
`MQLONG PutMsgRecFields=MQPMRF_MSG_ID + MQPMRF_CORREL_ID;`

Next, the sample allocates the response records and object records. The object records (MQORs) require at least one pair of names and an even number of names, that is, *ObjectName* and *ObjectQMGrName*.

The next stage involves connecting to the queue managers using MQCONN. The sample attempts to connect to the queue manager associated with the first queue in the MQOR; if this fails, it goes through the object records in turn. You are informed if it is not possible to connect to any queue manager and the program exits.

The target queues are opened using MQOPEN and the message is put to these queues using MQPUT. Any problems and failures are reported in the response records (MQRRs).

Finally, the target queues are closed using MQCLOSE and the program disconnects from the queue manager using MQDISC. The same response records are used for each call stating the *CompCode* and *Reason*.

The Echo sample programs:

The Echo sample programs echo a message from a message queue to the reply queue.

See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

The programs are intended to run as triggered programs.

On IBM i, UNIX, Linux, and Windows systems, their only input is an MQTMC2 (trigger message) structure that contains the name of a target queue and the queue manager. The COBOL version uses the default queue manager.

IBM i On IBM i, for the triggering process to work, ensure that the Echo sample program that you want to use is triggered by messages arriving on queue SYSTEM.SAMPLE.ECHO. To do this, specify the name of the Echo sample program that you want to use in the *ApplId* field of the process definition SYSTEM.SAMPLE.ECHOPROCESS. (For this, you can use the CHGMQMPCRC command; for details, see Change MQ Process (CHGMQMPCRC).) The sample queue has a trigger type of FIRST, so, if there are already messages on the queue before you run the Request sample, the Echo sample is not triggered by the messages that you send.

When you have set the definition correctly, first start AMQSERV4 in one job, then start AMQSREQ4 in another. You could use AMQSTRG4 instead of AMQSERV4, but potential job submission delays could make it less easy to follow what is happening.

Use the Request sample programs to send messages to queue SYSTEM.SAMPLE.ECHO. The Echo sample programs send a reply message containing the data in the request message to the reply-to queue specified in the request message.

Design of the Echo sample programs

The program opens the queue named in the trigger message structure that it was passed when it started. (For clarity, this is referred to as the request queue.) The program uses the MQOPEN call to open this queue for shared input.

The program uses the MQGET call to remove messages from this queue. This call uses the MQGMO_ACCEPT_TRUNCATED_MSG, MQGMO_CONVERT, and MQGMO_WAIT options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message and displays a warning message.

For each line of input, the program then reads the text into a buffer and uses the MQPUT1 call to put a request message, containing the text of that line, onto the reply-to queue.

If the MQGET call fails, the program puts a report message on the reply-to queue, setting the *Feedback* field of the message descriptor to the reason code returned by the MQGET.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

IBM i On IBM i, the program can also respond to messages sent to the queue from platforms other than IBM MQ for IBM i, although no sample is supplied for this situation. To make the ECHO program work:

- Write a program, correctly specifying the **Format**, **Encoding**, and **CCSID** parameters, to send text request messages.

The ECHO program requests the queue manager to perform message data conversion, if this is needed.

- Specify CONVERT(*YES) on the IBM MQ for IBM i sending channel, if the program that you have written does not provide similar conversion for the reply.

The Get sample programs:

The Get sample programs get messages from a queue using the MQGET call.

See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

Design of the Get sample program

The program opens the target queue using the MQOPEN call with the MQOO_INPUT_AS_Q_DEF option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

For each message on the queue, the program uses the MQGET call to remove the message from the queue, then displays the data contained in the message. The MQGET call uses the MQGMO_WAIT option, specifying a *WaitInterval* of 15 seconds, so that the program waits for this period if there is no message on the queue. If no message arrives before this interval expires, the call fails and returns the MQRC_NO_MSG_AVAILABLE reason code.

The program demonstrates how you must clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The MQGET call specifies a buffer of fixed size. If a message is longer than this buffer, the call fails and the program stops.

The program continues until either the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code or the MQGET call fails. If the call fails, the program displays an error message that contains the reason code.

The program then closes the queue using the MQCLOSE call.

Running the amqsget and amqsgetc samples:

These programs each take the following positional parameters:

1. The name of the source queue (required)
2. The name of the queue manager (optional)
If a queue manager is not specified, amqsget connects to the default queue manager and amqsgetc connects to the queue manager identified by an environment variable or the client channel definition file.
3. The open options (optional)
If open options are not specified, the sample uses a value of 8193 which is the combination of these two options:
 - MQOO_INPUT_AS_Q_DEF
 - MQOO_FAIL_IF QUIESCING
4. The close options (optional)
If close options are not specified, the sample uses a value of 0 which is MQCO_NONE.

These programs also use an environment variable named **MQSAMP_USER_ID** which should be set to the user ID to be used for connection authentication. When this is set, the program will prompt for a password to accompany that user ID.

To run these programs, enter one of the following:

- `amqsget myqueue qmanagername`
- `amqsgetc myqueue qmanagername`

where `myqueue` is the name of the queue from which the program will get messages, and `qmanagername` is the queue manager that owns `myqueue`.

High availability sample programs:

The **amqsgfhac**, **amqsphac**, and **amqsmfhac** high availability sample programs use automated client reconnection to demonstrate recovery following the failure of a queue manager. **amqsfhac** checks that a queue manager using networked storage maintains data integrity following a failure.

The **amqsgfhac**, **amqsphac**, and **amqsmfhac** programs are started from the command line, and can be used in combination to demonstrate reconnection after the failure of one instance of a multi-instance queue manager.

Alternatively, you can also use the **amqsgfhac**, **amqsphac**, and **amqsmfhac** samples to demonstrate client reconnection to single instance queue managers, typically configured into a queue manager group.

To keep the example simple, so it is easy to configure, you are shown the sample programs reconnecting to a single instance queue manager that is started, stopped and then restarted again; see “Set up and control the queue manager” on page 495.

Use **amqsfhac** in parallel with **amqmfscck** to check file system integrity. See **amqmfscck** (file system check) and Verifying shared file system behavior for more information.

amqsphac *queueName [qMgrName]*

- **amqsphac** is an IBM MQ MQI client application. It puts a sequence of messages to a queue with a two second delay between each message and displays events sent to its event handler.
- No sync point is used to put messages to the queue.
- Reconnection can be made to any queue manager in the same queue manager group.

amqsgfhac *queueName [qMgrName]*

- **amqsgfhac** is an IBM MQ MQI client application. It gets messages from a queue and displays events sent to its event handler.
- No sync point is used to get messages from the queue.
- Reconnection can be made to any queue manager in the same queue manager group.

amqsmfhac *-s sourceQueueName -t targetQueueName [-m qMgrName] [-w waitInterval]*

- **amqsmfhac** is an IBM MQ MQI client application. It copies messages from one queue to another with a default wait interval of 15 minutes after the last message that is received before the program finishes.
- The messages are copied within sync point.
- Reconnection can be made only to the same queue manager.

amqsfhac *QueueManagerName QueueName SideQueueName InTransactionCount RepeatCount (0 | 1 | 2)*

- **amqsfhac** is an IBM MQ MQI client application. It checks that an IBM MQ multi-instance queue manager using networked storage, such as a NAS or a cluster file system, maintains data integrity. Follow the steps to run **amqsfhac** in Verifying shared file system behavior.

- It uses the MQCNO_RECONNECT_Q_MGR option when connecting to *QueueManagerName*. It automatically reconnects when the queue manager fails over.
- It puts *InTransactionCount*RepeatCount* persistent messages to *QueueName* during which time you cause the queue manager to fail over any number of times. **amqsfhac** reconnects to the queue manager each time, and continues. The test is to make sure that no messages are lost.
- *InTransactionCount* messages are put within each transaction. The transaction is repeated *RepeatCount* number of times. If a failure occurs within a transaction, **amqsfhac** rolls back and resubmits the transaction when **amqsfhac** reconnects to the queue manager.
- It also puts messages to *SideQueueName*. It uses *SideQueueName* to check whether the all the messages are committed or rolled back from *QueueName* successfully. If it detects an inconsistency, it writes out an error message.
- Vary the amount of output tracing from **amqsfhac** by setting the last parameter to (0|1|2).

0	Least output.
1	Middling output.
2	Most output.

Configuring a client connection

You need to configure a client and server connection channel to run the samples. The client verification procedure explains how to set up a client test environment.

Alternatively, use the configuration provided in the following example.

Example using amqsgfhac, amqsfhac, and amqsmfhac

The example demonstrates reconnectable clients using a single instance queue manager.

Messages are placed on the queue SOURCE by **amqsfhac**, transferred to TARGET by **amqsmfhac**, and retrieved from TARGET by **amqsgfhac** ; see Figure 84.




Figure 84. Reconnectable client samples

Follow these steps to run the samples.

1. Create a file `hasamples.tst` containing the commands:
2. Type the following commands at a command prompt:
 - a. `crtmqm QM1`
 - b. `strmqm QM1`
 - c. `runmqsc QM1 < hasamples.tst`
3. Set the environment variable **MQCHLLIB** to the path to the AMQCLCHL.TAB client channel definition file; for example, `SET MQCHLLIB=C:\IBM\MQ\MQ7\Data\qmgrs\QM1\@ipcc.`

4. Open three new windows with **MQCHLLIB** set; for example on Windows, type **start** three times at the previous command prompt starting each program in one of the windows. See step 5 in “Set up and control the queue manager.”)
5. Type the command `endmqm -r -p QM1` to stop the queue manager, and then allow the clients to reconnect.
6. Type the command `strmqm QM1` to restart the queue manager.

The results from running the **amqsghac**, **amqsphac**, and **amqsmhac** samples on Windows are shown in the following examples.

Set up and control the queue manager

1. Create the queue manager.

```
C:\> crtmqm QM1
IBM MQ queue manager created.
Directory 'C:\IBM\MQ\MQ7\Data\qmgrs\QM1' created.
Creating or replacing default objects for QM1.
Default objects statistics : 67 created. 0 replaced. 0 failed.
Completing setup.
Setup completed.
```

Remember the data directory to set the **MQCHLLIB** variable later.

2. Start the queue manager.

```
C:\> strmqm QM1

IBM MQ queue manager 'QM1' starting.
5 log records accessed on queue manager 'QM1' during the log replay phase.
Log replay for queue manager 'QM1' complete.
Transaction manager state recovered for queue manager 'QM1'.
IBM MQ queue manager 'QM1' started.
```

3. Create the queues and channels, modify the listener port, and start the listener and channel.
4. Make the client channel table known to the clients.

Use the data directory returned from the **crtmqm** command in step 1, and add the directory @ipcc to it to set the **MQCHLLIB** variable.

```
C:\> SET MQCHLLIB=C:\IBM\MQ\MQ7\Data\qmgrs\QM1\@ipcc
```

5. Start the sample programs in the other windows

```
C:\> start amqsphac SOURCE QM1
C:\> start amqsmhac -s SOURCE -t TARGET -m QM1
C:\> start amqsghac TARGET QM1
```

6. End the queue manager and restart it again.

```
C:\> endmqm -r -p QM1

Waiting for queue manager 'QM1' to end.
IBM MQ queue manager 'QM1' ending.
IBM MQ queue manager 'QM1' ended.

C:\> strmqm QM1

IBM MQ queue manager 'QM1' starting.
5 log records accessed on queue manager 'QM1' during the log replay phase.
Log replay for queue manager 'QM1' complete.
Transaction manager state recovered for queue manager 'QM1'.
IBM MQ queue manager 'QM1' started.
```

amqsphac

```
Sample AMQSPHAC start
target queue is SOURCE
message Message 1
message Message 2
```

```
16:25:22 : EVENT : Connection Reconnecting (Delay: 0ms)
16:25:45 : EVENT : Connection Reconnecting (Delay: 0ms)
16:26:02 : EVENT : Connection Reconnectedmessage
Message 3
message Message 4
message Message 5
```

amqsmhac

```
Sample AMQSMHA0 start
16:25:22 : EVENT : Connection Reconnecting (Delay: 0ms)
16:25:45 : EVENT : Connection Reconnecting (Delay: 0ms)
16:26:02 : EVENT : Connection Reconnected
No more messages.
Sample AMQSMHA0 end
C:\>
```

amqsgnac

```
Sample AMQSGHAC start
message Message 1
message Message 2
16:25:22 : EVENT : Connection Reconnecting (Delay: 0ms)
16:25:45 : EVENT : Connection Reconnecting (Delay: 0ms)
16:26:02 : EVENT : Connection Reconnected
message Message 3
message Message 4
message Message 5
```

Related information:

Verifying shared file system behavior

amqmfsc (file system check)


The Inquire sample programs:

The Inquire sample programs inquire about some of the attributes of a queue using the MQINQ call.



See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

These programs are intended to run as triggered programs, so their only input is an MQTMC2 (trigger message) structure for IBM i, Windows, UNIX and Linux systems. This structure contains the name of a target queue with attributes that are to be inquired upon. The C version also uses the queue manager name. The COBOL version uses the default queue manager.


For the triggering process to work, ensure that the Inquire sample program that you want to use is triggered by messages arriving on queue SYSTEM.SAMPLE.INQ. To do this, specify the name of the Inquire sample program that you want to use in the *ApplicId* field of the process definition

SYSTEM.SAMPLE.INQPROCESS.  For IBM i, you can use the CHGMQMPCR command for this; for details, see Change MQ Process (CHGMQMPCR). The sample queue has a trigger type of FIRST; if there are already messages on the queue before you run the request sample, the inquire sample is not triggered by the messages that you send.

When you have set the definition correctly:

-  For UNIX, Linux, and Windows, start the **runmqtrm** program in one session, then start the **amqsreq** program in another.
-  For IBM i, start the AMQSERV4 program in one session, then start the AMQSREQ4 program in another. You could use AMQSTRG4 instead of AMQSERV4, but potential job submission delays could make it less easy to follow what is happening.

Use the Request sample programs to send request messages, each containing just a queue name, to queue SYSTEM.SAMPLE.INQ. For each request message, the Inquire sample programs send a reply message containing information about the queue specified in the request message. The replies are sent to the reply-to queue specified in the request message.

 On IBM i, if the sample input file member QMQMSAMP.AMQSDATA(INQ) is used, the last queue named does not exist, so the sample returns a report message with a reason code for the failure.

Design of the Inquire sample program

The program opens the queue named in the trigger message structure that it was passed when it started. (For clarity, we will call this the *request queue*.) The program uses the MQOPEN call to open this queue for shared input.

The program uses the MQGET call to remove messages from this queue. This call uses the MQGMO_ACCEPT_TRUNCATED_MSG and MQGMO_WAIT options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message and displays a warning message.

For each request message removed from the request queue, the program reads the name of the queue (which we will call the *target queue*) contained in the data, and opens that queue using the MQOPEN call with the MQOO_INQ option. The program then uses the MQINQ call to inquire about the values of the *InhibitGet*, *CurrentQDepth*, and *OpenInputCount* attributes of the target queue.

If the MQINQ call is successful, the program uses the MQPUT1 call to put a reply message on the reply-to queue. This message contains the values of the three attributes.

If the MQOPEN or MQINQ call is unsuccessful, the program uses the MQPUT1 call to put a report message on the reply-to queue. In the *Feedback* field of the message descriptor of this report message is the reason code returned by either the MQOPEN or MQINQ call, depending on which one failed.

After the MQINQ call, the program closes the target queue using the MQCLOSE call.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

The Inquire Properties of a Message Handle sample program:

AMQSIQMA is a sample C program to inquire properties of a message handle from a message queue, and is an example of the use of the MQINQM API call.

This sample creates a message handle and puts it into the MsgHandle field of the MQGMO structure. The sample then gets one message and inquires and prints all properties with which the message handle was populated.

```
C:\Program Files\IBM\MQ\tools\c\Samples\Bin >amqsism Q QM1
Sample AMQSIQMA start
property name MyProp value MyValue
message text Hello world!
Sample AMQSIQMA end
```

The Publish/Subscribe sample programs:

The publish/subscribe sample programs demonstrate the use of the publish and subscribe features in IBM MQ.

There are three C language sample programs illustrating how to program to the IBM MQ publish/subscribe interface. There are some C samples that use older interfaces, and there are Java samples. The Java samples use the IBM MQ publish/subscribe interface in `com.ibm.mq.jar` and the JMS publish/subscribe interface in `com.ibm.mqjms`. The JMS samples are not covered in this topic.

C

Find the publisher sample `amqspub` in the C samples folder. Run it with any topic name you like as the first parameter, followed by an optional queue manager name. For example, `amqspub mytopic QM3`. There is also a client version called `amqspubc`. If you choose to run the client version, first see "Configuring a queue manager to accept client connections on Windows, UNIX and IBM i" on page 470 for details.

The publisher connects to the default queue manager and responds with the output, `target topic is mytopic`. Every line you enter into this window from now onwards is published to `mytopic`.

Open another command window in the same directory, and run the subscriber program, `amqssub`, supplying it with the same topic name, and an optional queue manager name. For example, `amqssub mytopic QM3`.

The subscriber responds with the output, `Calling MQGET : 30 seconds wait time`. From now onwards, lines you type into the publisher appear in the output of the subscriber.

Start another subscriber in another command window, and watch both subscribers receive publications.

For full documentation of the parameters, including setting options, refer to the sample source code. The values for the subscriber options field is described in the following topic: [Options \(MQLONG\)](#).

There is another subscriber sample `amqssbx`, which offers additional subscription options as command line switches.

Type `amqssbx -d mysub -t mytopic -k` to invoke the subscriber using durable subscriptions that are retained after the subscriber has terminated.

Test the subscription by publishing another item using the publisher. Wait for 30 seconds for the subscriber to terminate. Publish some more items under the same topic. Restart the subscriber. The last item published while the subscriber was not running is displayed by the subscriber immediately it is restarted.

C legacy

There is an additional set of C samples which demonstrate queued commands. Some of these samples were originally shipped as part of the MQ0C Supportpac. The capabilities the samples demonstrate are fully supported, for compatibility reasons.

We discourage you from using the queued command interface. It is much more complex than the publish/subscribe API, and there is no compelling functional reason to program complex queued commands. However, you might find the queued approach more suitable, perhaps because you are already using the interface, or because your programming environment makes it easier to build a complex message and call a generic `MQPUT`, rather than constructing different calls to `MQSUB`.

The additional samples are located in the `pubsub` subdirectory in the `samples` folder.

There are six types of sample listed in Table 76.

Table 76. Categories of legacy publish/subscribe sample C programs

Category	Programs	Comments
RFH1	amqssr1a.c amqspr1a.c	Simple publish/subscribe example built using RFH1 format messages.
RFH2	amqssr2a.c amqssp2a.c	Simple publish/subscribe example built using RFH2 format messages.
MQAI samples	amqspca.c amqsspca.c	Simple publish/subscribe example built using PCF commands and the MQAI command interface.
MA0C Results service using RFH1	amqsgama.c amqsresa.c	Results service built using RFH1 headers 1. Requires the queues defined in amqsgama.tst and amqsresa.tst 2. amqsresa must be started before amqsgama
MA0C Results service using RFH2	amqmgr2a.c amqrr2a.c	Results service built using RFH2 headers 1. Requires the queues defined in amqsgama.tst and amqsresa.tst 2. amqsresa must be started before amqsgama
Routing exit publish/subscribe sample	amqspdra.c	Demonstrates how to change the queue or queue manager destination for a publish/subscribe message in a routing exit.

Java

The Java sample `MQPubSubApiSample.java` combines publisher and subscribers in a single program. Its source and compiled class files are found in the `wmqjava` samples folder.

If you choose to run in client mode, first see “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470 for details.

Run the sample from the command line using the Java command, if you have a Java environment configured. You can also run the sample from the IBM MQ Explorer Eclipse workspace that has a Java programming workbench already set up.

You might need to change some of the sample program's properties to run it. You do this by providing parameters to the JVM, or editing the source.

The instructions in “Running the MQPubSubApiSample Java sample” on page 500 show how to run the sample from the Eclipse workspace.

Running the MQPubSubApiSample Java sample:

How to run the MQPubSubApiSample using the Java Development Tools from the Eclipse platform.

Before you begin

Open the Eclipse workbench. Create a new workspace directory and select it. Close the welcome window.

Follow the steps in “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470 prior to running as a client.

About this task

The Java publish/subscribe sample program is an IBM MQ MQI client Java program. The sample runs without modification using a default queue manager listening on port 1414. The task describes this simple case, and indicates in general terms how to provide parameters and modify the sample to suit different IBM MQ configurations. The example is illustrated running on Windows. The file paths will differ on other platforms.

Procedure

1. Import the Java sample programs
 - a. In the workbench, click **Window** > **Open perspective** > **Other** > **Java** and click **OK**.
 - b. Switch to the Package Explorer view.
 - c. Right-click in the white-space in the Package Explorer view. Click **New** > **Java project**.
 - d. In the **Project name** field type MQ Java Samples. Click **Next**.
 - e. In the **Java Settings** panel, switch to the Libraries tab.
 - f. Click **Add External JARs**.
 - g. Browse to *MQ_INSTALLATION_PATH* \java\lib where *MQ_INSTALLATION_PATH* is the IBM MQ installation folder and select com.ibm.mq.jar and com.ibm.mq.jmqi.jar
 - h. Click **Open** > **Finish**.
 - i. Right-click src in the Package Explorer view.
 - j. Select **Import...** > **General** > **File System** > **Next** > **Browse...** and browse to the path *MQ_INSTALLATION_PATH* \tools\wmqjava\samples where *MQ_INSTALLATION_PATH* is the IBM MQ installation directory.
 - k. On the Import panel, Figure 85 on page 501, click samples (do not select the check box).
 - l. Select MQPubSubApiSample.java. The **Into folder** field should contain MQ Java Samples/src. Click **Finish**.

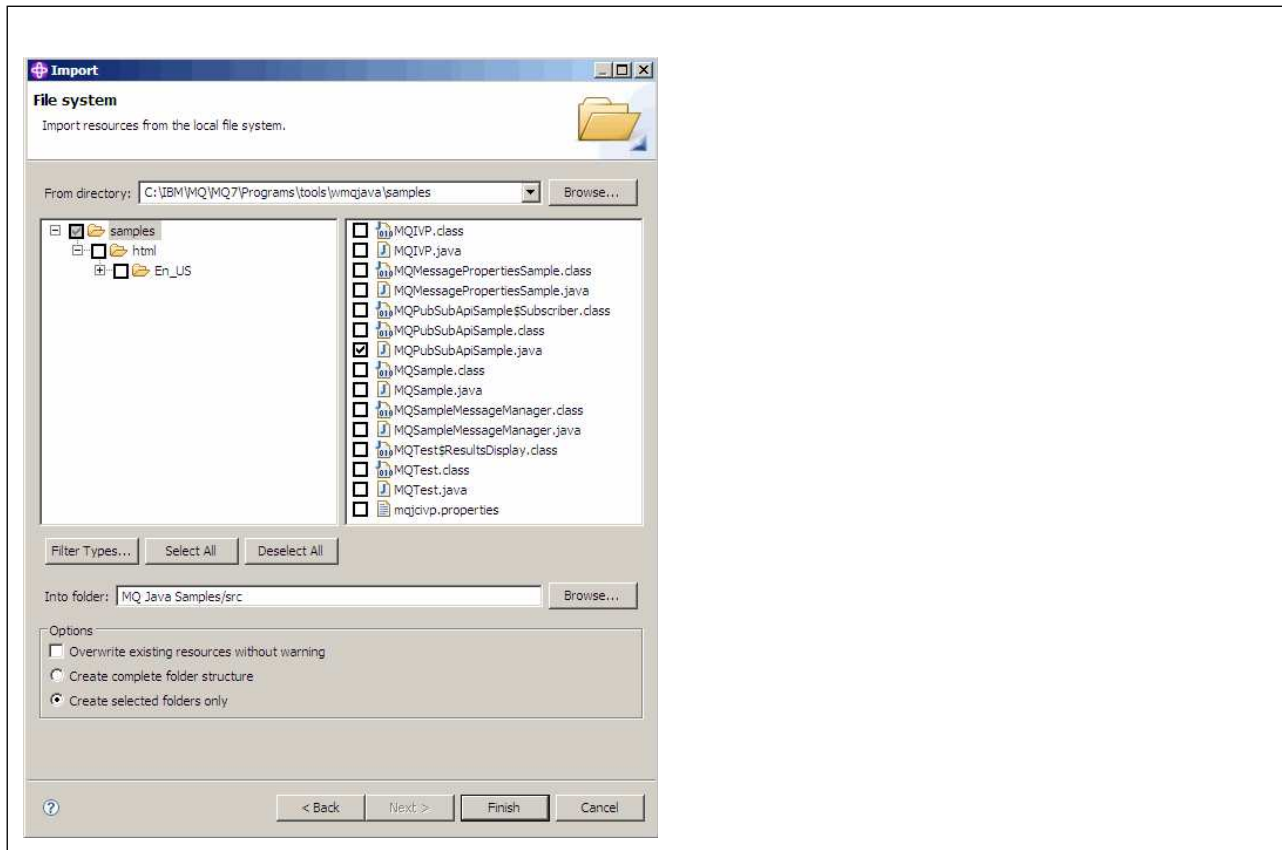


Figure 85. File system import

2. Run the publish/subscribe sample program. There are two ways to run the program, depending on whether you need to change the default parameters.

- The first choice runs the program without making any changes:
 - In the workspace main menu, expand the src folder. Right-click **MQPubSubApiSample.java**
Run-as > 1. Java Application
- The second choice runs the program with parameters or with modified source code for your environment:

- Open MQPubSubApiSample.java and study the MQPubSubApiSample constructor.
- Modify the attributes of the program.

These attributes are modifiable using the `-D JVM` switch, or by providing a default value for the System property by editing the source code.

- topicObject
- queueManagerName
- subscriberCount

These attributes are changeable only by editing the source code in the constructor.

- hostname
- port
- channel

To set System properties, code a default value in the accessor, for example:

```
queueManagerName = System.getProperty("com.ibm.mq.pubSubSample.queueManagerName", "QM3");
```

Or provide the parameter to the JVM using the `-D` option, as shown in the following steps:

- a. Copy the full name of the System.Property you want to set, for example:
`com.ibm.mq.pubSubSample.queueManagerName.`



- b. In the workspace, right-click **Run > Open Run Dialog**. Double-click Java Application in Create, Manage and Run applications and click the (x) = Arguments tab.
- c. In the VM arguments: pane, type `-D` and paste the System.property name, `com.ibm.mq.pubSubSample.queueManagerName`, followed by `=QM3`. Click **Apply > Run**.
- d. Add further arguments as a comma separated list, or as additional lines in the pane, without comma separators.
For example: `-Dcom.ibm.mq.pubSubSample.queueManagerName=QM3,`
`-Dcom.ibm.mq.pubSubSample.subscriberCount=6.`

The Publish Exit sample program:

AMQSPSE0 is a sample C program of an exit to intercept a publication before it is delivered to a subscriber. The exit can then, for example, alter the message headers, payload or destination, or prevent the message being published to a subscriber.


To run the sample, perform the following tasks:



1. Configure the queue manager:


-   On UNIX and Linux systems add a stanza like this to the `qm.ini` file:

```
PublishSubscribe:
PublishExitPath=Module
PublishExitFunction=EntryPoint
```

where the module is `MQ_INSTALLATION_PATH/samp/bin/amqspse`. `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

-  On Windows set the equivalent attributes in the registry.
2. Make sure the Module is accessible to IBM MQ.
 3. Restart the queue manager to pick up the configuration.
 4. In the application process to be traced, describe where the trace files should be written to. For example:

-   On UNIX and Linux systems, ensure that the directory `/var/mqm/trace` exists and export the following environment variable:

```
export MQPSE_TRACE_LOGFILE=/var/mqm/trace/PubTrace
```
-  On Windows, ensure that the directory `C:\temp` exists and set the following environment variable:

```
set MQPSE_TRACE_LOGFILE=C:\temp\PubTrace
```

The Put sample programs:

The Put sample programs put messages on a queue using the MQPUT call.

See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

Design of the Put sample program

The program uses the MQOPEN call with the MQOO_OUTPUT option to open the target queue for putting messages.

If it cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call. To keep the program simple, on this and on subsequent MQI calls, the program uses default values for many of the options.

For each line of input, the program reads the text into a buffer and uses the MQPUT call to create a datagram message containing the text of that line. The program continues until either it reaches the end of the input or the MQPUT call fails. If the program reaches the end of the input, it closes the queue using the MQCLOSE call.

Running the Put sample programs:



Running the amqsput and amqsputc samples

The amqsput sample is the program for putting messages using local bindings, and the amqsputc sample is the program for putting using client bindings. These programs each take the following positional parameters:

1. The name of the target queue (required)
2. The name of the queue manager (optional)
If a queue manager is not specified, amqsput connects to the default queue manager and amqsputc connects to the queue manager identified by the MQSERVER environment variable or the client channel definition file.
3. The open options (optional)
If open options are not specified, the sample uses a value of 8208 which is the combination of these two options:
 - MQOO_OUTPUT
 - MQOO_FAIL_IF QUIESCING
4. The close options (optional)
If close options are not specified, the sample uses a value of 0 which is MQCO_NONE.
5. The name of the target queue manager (optional)
If a target queue manager is not specified, the ObjectQMgrName field in the MQOD will be left blank.
6. The name of the dynamic queue (optional)
If a dynamic queue name is not specified, the DynamicQName field in the MQOD will be left blank.

These programs also use an environment variable named **MQSAMP_USER_ID** which should be set to the user ID to be used for connection authentication. When this is set, the program will prompt for a password to accompany that user ID.

To run these programs, enter one of the following:

- amqsput myqueue qmanagername

- `amqsputc myqueue qmanagername`

where `myqueue` is the name of the queue on which the messages are going to be put, and `qmanagername` is the queue manager that owns `myqueue`.



Running the `amq0put` sample

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

Please enter the name of the target queue

It takes input from StdIn and adds each line of input to the target queue. A blank line indicates there is no more data.



Running the `AMQSPUT4 C` sample (IBM i)

The C program `AMQSPUT4`, available only for the IBM i platform, creates messages by reading data from a member of a source file.

You must specify the name of the file as a parameter when you start the program. The structure of the file must be:

```
queue name
text of message 1
text of message 2
:
text of message n
blank line
```

A sample of input for the put samples is supplied in library `QMOMSAMP` file `AMQSDATA` member `PUT`.

Note: Remember that queue names are case sensitive. All the queues created by the sample file create program `AMQSAMP4` have names created in uppercase characters.

The C program puts messages on the queue named in the first line of the file; you can use the supplied queue `SYSTEM.SAMPLE.LOCAL`. The program puts the text of each of the following lines of the file into separate datagram messages, and stops when it reads a blank line at the end of the file.

Using the example data file the command is:

```
CALL PGM(QMQM/AMQSPUT4) PARM('QMOMSAMP/AMQSDATA(PUT)')
```



Running the `AMQ0PUT4 COBOL` sample (IBM i)

The COBOL program `AMQ0PUT4`, available only on the IBM i platform, creates messages by accepting data from the keyboard.


To start the program, call the program and give the name of your target queue as a program parameter. The program accepts input from the keyboard into a buffer and creates a datagram message for each line of text. The program stops when you enter a blank line at the keyboard.


The Reference Message sample programs:

The Reference Message samples allow a large object to be transferred from one node to another (usually on different systems) without the need for the object to be stored on IBM MQ queues at either the source or the destination nodes.

A set of sample programs is provided to demonstrate how Reference Messages can be put to a queue, received by message exits, and taken from a queue. The sample programs use Reference Messages to move files. If you want to move other objects such as databases, or if you want to perform security checks, define your own exit, based on the sample, amqsxrm.

The version of the Reference Message exit sample program to use depends on the platform on which the channel is running:

- On all platforms, use amqsxrma at the sending end.
- Use amqsxrma at the receiving end if the receiver is running under any platform except IBM i.
-  If the receiver is running under IBM i, use amqsxrm4.

Notes for IBM i users: 

To receive a Reference Message using the sample message exit, specify a file in the root file system of IFS or any subdirectory so that a stream file can be created.

The sample message exit on IBM i creates the file, converts the data to EBCDIC, and sets the code page to your system code page. You can then copy this file to the QSYS.LIB file system using the CPYFRMSTMF command. For example:

```
CPYFRMSTMF FROMSTMF('JANEP/TEST.TXT')
TOMBR('qsys.lib.janep.lib/test.fie/test.mbr') MBROPT(*REPLACE)
CVTDTA(*NONE)
```

The CPYFRMSTMF command does not create the file. You must create it before running this command.

If you send a file from QSYS.LIB, no changes are required to the samples. For any other file system ensure that the CCSID specified in the CodedCharSetId field in the MQRMH structure matches the bulk data that you are sending.

When using the integrated file system, create program modules with the SYSIFCOPT(*IFSIO) option set. If you want to move database or fixed-length record files, define your own exit based on the supplied sample AMQSXRM4.

The recommended method of transferring a database file is to convert it to IFS structure, using the CPYTOSTMF command, and then send the Reference Message attaching the IFS file. If you choose to transfer a database file by referring to it from within IFS, but do not convert it to IFS structure, you must specify the member name. Data integrity is not guaranteed if you choose this method.

The application that receives the Reference Message (AMQSGRM or AMQSGRMA on IBM i) is triggered using the initiation queue INITQ and process PROC. Ensure that the CONNAME fields are set correctly and the MSGEXIT field reflects your directory structure, depending on machine type and where the IBM MQ product is installed.

The MQSC definitions have used an AIX style for defining the exits, so if you are using MQSC on IBM i, you need to modify these accordingly. It is important to note that the message data FLATFILE is case sensitive and the sample will not work unless it is in uppercase.

On machine MACHINE1, queue manager QMGR1

MQSC syntax

```
define chl(chl1) chltype(sdr) trptype(tcp) conname('machine2') xmitq(xmitq)
msgdata(FLATFILE) msgexit('/usr/lpp/mqm/samp/bin/amqsxrm(MsgExit)
')
```

```
define ql(xmitq) usage(xmitq)
```

```
define chl(report) chltype(rcvr) trptype(tcp) replace
```

```
define qr(qr) rname(dq) rqmname(qmgr2) xmitq(xmitq) replace
```

IBM i command syntax

Note: If you do not specify a queue manager name the system uses the default queue manager.

```
CRTMQMCHL CHLNAME(CHL1) CHLTYPE(*SDR) MQMNAME(QMGR1) +
REPLACE(*YES) TRPTYPE(*TCP) +
CONNAME('MACHINE2(60501)') TMQNAME(XMITQ) +
MSGEXIT(QMQM/AMQSXRMA) MSGUSRDATA(FLATFILE)
```

```
CRTMQMQ QNAME(XMITQ) QTYPE(*LCL) MQMNAME(QMGR1) +
REPLACE(*YES) USAGE(*TMQ)
```

```
CRTMQMCHL CHLNAME(REPORT) CHLTYPE(*RCVR) +
MQMNAME(QMGR1) REPLACE(*YES) TRPTYPE(*TCP)
```

```
CRTMQMQ QNAME(QR) QTYPE(*RMT) MQMNAME(QMGR1) +
REPLACE(*YES) RMTQNAME(DQ) +
RMTMQMNAME(QMGR2) TMQNAME(XMITQ)
```

On machine MACHINE2, queue manager QMGR2

MQSC syntax

```
define chl(chl1) chltype(rcvr) trptype(tcp)
msgexit('/usr/lpp/mqm/samp/bin/amqsxrm(MsgExit)')
msgdata(flatfile)
```

```
define chl(report) chltype(sdr) trptype(tcp) conname('MACHINE1')
xmitq(qmgr1)
```


```
define ql(initq)
```

```
define ql(qmgr1) usage(xmitq)
```

```
define pro(proc) applid('/usr/lpp/mqm/samp/bin/amqsgrm')
```

```
define ql(dq) initq(initq) process(proc) trigger trigtype(first)
```

IBM i command syntax

Note:  If you do not specify a queue manager name the system uses the default queue manager.

```
CRTMQMCHL CHLNAME(CHL1) CHLTYPE(*RCVR) MQMNAME(QMGR2) +
REPLACE(*YES) TRPTYPE(*TCP) +
MSGEXIT(QMQM/AMQSXRMA) MSGUSRDATA(FLATFILE)
```

```
CRTMQMCHL CHLNAME(REPORT) CHLTYPE(*SDR) MQMNAME(QMGR2) +  
REPLACE(*YES) TRPTYPE(*TCP) +  
CONNNAME('MACHINE1(60500)') TMQNAME(QMGR1)
```

```
CRTMQMQ QNAME(INITQ) QTYPE(*LCL) MQMNAME(QMGR2) +  
REPLACE(*YES) USAGE(*NORMAL)
```

```
CRTMQMQ QNAME(QMGR1) QTYPE(*LCL) MQMNAME(QMGR2) +  
REPLACE(*YES) USAGE(*TMQ)
```

```
CRTMQMPC PRCNAME(PROC) MQMNAME(QMGR2) REPLACE(*YES) +  
APPID('QMOM/AMQSGRM4')
```

```
CRTMQMQ QNAME(DQ) QTYPE(*LCL) MQMNAME(QMGR2) +  
REPLACE(*YES) PRCNAME(PROC) TRGENBL(*YES) +  
INITQNAME(INITQ)
```

2. Once the IBM MQ objects have been created:
 - a. Where applicable to the platform, start the listener for the sending and receiving queue managers
 - b. Start the channels CHL1 and REPORT
 - c. On the receiving queue manager start the trigger monitor for the initiation queue INITQ
3. Invoke the put Reference Message sample program AMQSPRM (AMQSPRMA on IBM i) from the command line using the following parameters:

- m Name of the local queue manager; this defaults to the default queue manager
- i Name and location of source file
- o Name and location of the destination file
- q Name of the queue
- g Name of the queue manager where the queue, defined in the -q parameter exists. This defaults to the queue manager specified in the -m parameter.
- t Object type
- w Wait interval, that is, the waiting time for exception and COA reports from the receiving queue manager

For example, to use the sample with the objects defined previously you would use the following parameters:

```
-mQMGR1 -iInput File -oOutput File -qQR -tFLATFILE -w120
```

Increasing the waiting time allows time for a large file to be sent across a network before the program putting the messages times out.

```
amqsprm -q QR -m QMGR1 -i d:\x\file.in -o d:\y\file.out -t FLATFILE
```


IBM i users:

- a. Use the following command:

```
CALL PGM(QMOM/AMQSPRM4) PARM('-mQMGR1' +  
'-i/refmsgs/rmsg1' +  
'-o/refmsgs/rmsgx' '-qQR' +  
'-gQMGR1' '-tFLATFILE' '-w15')
```

This assumes that the original file rmsg1 is in IFS directory /refmsgs and that you want the destination file to be rmsgx in IFS directory /refmsgs on the target system.

- b. Create your own directory using the CRTDIR command rather than using the root directory.
- c. When you call the program that puts data, remember that the output file name needs to reflect the IFS naming convention; for example /TEST/FILENAME creates a file called FILENAME in the directory TEST.

Note:  You can use either a forward slash (/) or a dash (-) when specifying parameters.

 For example:

```
amqsprmq /i d:\files\infile.dat /o e:\files\outfile.dat /q QR  
/m QMGR1 /w 30 /t FLATFILE
```

Note: For UNIX and Linux platforms, you must use two backslashes (\\) instead of one to denote the destination file directory. Therefore, the **amqsprmq** command looks like this:




```
amqsprmq -i /files/infile.dat -o e:\\files\\outfile.dat -q QR  
-m QMGR1 -w 30 -t FLATFILE
```

Running the put Reference Message program does the following:

- The Reference Message is put to queue QR on queue manager QMGR1.
 - The source file and path are d:\files\infile.dat and exists on the system where the example command is issued.
 - If the queue QR is a remote queue, the Reference Message is sent to another queue manager, on a different system, where a file is created with the name and path e:\files\outfile.dat. The contents of this file are the same as the source file.
 - amqsprmq waits for 30 seconds for a COA report from the destination queue manager.
 - The object type is flatfile, so the channel used to move messages from the queue QR must specify this in the *MsgData* field.
4. When you define your channels, select the message exit at both the sending and receiving ends to be amqsxrm.

 This is defined on Windows as follows:




```
msgexit(' pathname\amqsxrm.dll(MsgExit)')
```

   This is defined on AIX, HP-UX, and Solaris as follows:

```
msgexit(' pathname/amqsxrm(MsgExit)')
```

If you specify a path name, specify the complete name. If you omit the path name, it is assumed that the program is in the path specified in the *qm.ini* file (or, on IBM MQ for Windows, the path specified in the registry).

5. The channel exit reads the Reference Message header and finds the file that it refers to.
6. The channel exit can then segment the file before sending it down the channel along with the header.

   On AIX, HP-UX, and Solaris, change the group owner of the target directory to 'mqm' so that the sample message exit can create the file in that directory. Also, change the permissions of the target directory to allow mqm group members to write to it. The file data is not stored on the IBM MQ queues.

7. When the last segment of the file is processed by the receiving message exit, the Reference Message is put to the destination queue specified by amqsprmq. If this queue is triggered (that is, the definition specifies **Trigger**, **InitQ**, and **Process** queue attributes), the program specified by the PROC parameter of the destination queue is triggered. The program to be triggered must be defined in the *AppId* field of the **Process** attribute.
8. When the Reference Message reaches the destination queue (DQ), a COA report is sent back to the putting application (amqsprmq).
9. The Get Reference Message sample, amqsgrm, gets messages from the queue specified in the input trigger message and checks the existence of the file.

Design of the Put Reference Message sample (amqsprma.c, AMQSPRM4):

This topic gives a detailed description of a Put Reference Message sample.

This sample creates a Reference Message that refers to a file and puts it on a specified queue:

1. The sample connects to a local queue manager using MQCONN.
2. It then opens (MQOPEN) a model queue that is used to receive report messages.
3. The sample builds a Reference Message containing the values required to move the file, for example, the source and destination file names and the object type. As an example, the sample shipped with IBM MQ builds a Reference Message to send the file `d:\x\file.in` from QMGR1 to QMGR2 and to re-create the file as `d:\y\file.out` using the following parameters:

```
amqsprm -q QR -m QMGR1 -i d:\x\file.in -o d:\y\file.out -t FLATFILE
```

Where QR is a remote queue definition that refers to a target queue on QMGR2.

Note: For UNIX and Linux platforms, use two backslashes (\\) instead of one to denote the destination file directory. Therefore, the `amqsprm` command looks like this:

```
amqsprm -q QR -m QMGR1 -i /x/file.in -o d:\\y\\file.out -t FLATFILE
```

4. The Reference Message is put (without any file data) to the queue specified by the `/q` parameter. If this is a remote queue, the message is put to the corresponding transmission queue.
5. The sample waits, for the duration of time specified in the `/w` parameter (which defaults to 15 seconds), for COA reports, which, along with exception reports, are sent back to the dynamic queue created on the local queue manager (QMGR1).

Design of the Reference Message Exit sample (amqsxrma.c, AMQSXRMA4):

This sample recognizes Reference Messages with an object type that matches the object type in the message exit user data field of the channel definition.

For these messages, the following happens:

- At the sender or server channel, the specified length of data is copied from the specified offset of the specified file into the space remaining in the agent buffer after the Reference Message. If the end of the file is not reached, the Reference Message is put back on the transmission queue after updating the `DataLogicalOffset` field.
- At the requester or receiver channel, if the `DataLogicalOffset` field is zero and the specified file does not exist, it is created. The data following the Reference Message is added to the end of the specified file. If the Reference Message is not the last one for the specified file, it is discarded. Otherwise, it is returned to the channel exit, without the appended data, to be put on the target queue.

For sender and server channels, if the `DataLogicalLength` field in the input Reference Message is zero, the remaining part of the file, from `DataLogicalOffset` to the end of the file, is to be sent along the channel. If it is not zero, only the length specified is sent.

If an error occurs (for example, if the sample cannot open a file), `MQCXP.ExitResponse` is set to `MQXCC_SUPPRESS_FUNCTION` so that the message being processed is put to the dead-letter queue instead of continuing to the destination queue. A feedback code is returned in `MQCXP.Feedback` and returned to the application that put the message in the `Feedback` field of the message descriptor of a report message. This is because the putting application requested exception reports by setting `MQRO_EXCEPTION` in the `Report` field of the `MQMD`.

If the encoding or `CodedCharacterSetId` (CCSID) of the Reference Message is different from that of the queue manager, the Reference Message is converted to the local encoding and CCSID. In our sample, `amqsprm`, the format of the object is `MQFMT_STRING`, so `amqsxrm` converts the object data to the local CCSID at the receiving end before the data is written to the file.

Do not specify the format of the file being transferred as MQFMT_STRING if the file contains multibyte characters (for example, DBCS or Unicode). This is because a multibyte character could be split when the file is segmented at the sending end. To transfer and convert such a file, specify the format as something other than MQFMT_STRING so that the Reference Message exit does not convert it and convert the file at the receiving end when the transfer is complete.

Compiling the Reference Message Exit sample:

To compile the Reference Message Exit sample, use the command for the platform on which IBM MQ is installed.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

To compile `amqsxrma`, use the following commands:

AIX

On AIX

```
xlc_r -q64 -e MsgExit -bE:amqsxrm.exp -bM:SRE -o amqsxrm_64_r  
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -lmqm_r amqsqrma.c
```

HP-UX

On HP-UX

```
$ cc89 +DD64 +z -c -D_HPUX_SOURCE -o amqsxrma.o amqsqrma.c -I MQ_INSTALLATION_PATH/inc  
$ ld -b amqsxrma.o -o /var/mqm/exits64/amqsxrma -L MQ_INSTALLATION_PATH/lib64  
-L/usr/lib/pa20_64 -lmqm_r -lpthread
```

IBM i

On IBM i

```
CRTCMOD MODULE(MYLIB/AMQXRMA) SRCFILE(QMQMSAMP/QCSRC)  
TERASPACE(*YES *TSIFC)
```

Note:

1. To create your module so that it uses the IFS file system, add the option `SYSIFCOPT(*IFSIO)`
2. To create the program for use with non-threaded channels use the following command: `CRTPGM PGM(MYLIB/AMQXRMA) BNDSRVPGM(QMQM/LIBMQM)`
3. To create the program for use with threaded channels use the following command: `CRTPGM PGM(MYLIB/AMQXRMA) BNDSRVPGM(QMQM/LIBMQM_R)`

Linux

On Linux

```
$ gcc -m64 -shared -fPIC -o /var/mqm/exits64/amqsxrma amqsqrma.c -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqm_r
```

Solaris

On Solaris

```
$ cc -xarch=v9 -mt -G -o /var/mqm/exits64/amqsxrma amqsqrma.c -I MQ_INSTALLATION_PATH/inc  
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqm
```

```
-lsocket  
-lnsl -ldl
```

Windows

On Windows

IBM MQ now supplies the mqm library with client packages as well as server packages, so the following example uses mqm.lib instead of mqmvx.lib:

```
cl amqsqrma.c /link /out:amqsxrm.dll /dll mqm.lib mqm.lib /def:amqsxrm.def
```

Related concepts:

“Writing channel-exit programs” on page 344

You can use the following information to help you write channel-exit programs.

Design of the Get Reference Message sample (amqsqrma.c, AMQSGRM4):

This topic explains the design of the Get Reference Message sample.

The program logic is as follows:

1. The sample is triggered and extracts the queue and queue manager names from the input trigger message.
2. It then connects to the specified queue manager using MQCONN and opens the specified queue using MQOPEN.
3. The sample issues MQGET with a wait interval of 15 seconds within a loop to get messages from the queue.
4. If a message is a Reference Message, the sample checks the existence of the file that has been transferred.
5. It then closes the queue and disconnects from the queue manager.

The Request sample programs:

The Request sample programs demonstrate client/server processing. The samples are the clients that put request messages on a target server queue that is processed by a server program. They wait for the server program to put a reply message on a reply-to queue.

The Request samples put a series of request messages on the target server queue using the MQPUT call. These messages specify the local queue, SYSTEM.SAMPLE.REPLY as the reply-to queue, which can be a local or remote queue. The programs wait for reply messages, then display them. Replies are sent only if the target server queue is being processed by a server application, or if an application is triggered for that purpose (the Inquire, Set, and Echo sample programs are designed to be triggered). The C sample waits 1 minute (the COBOL sample waits 5 minutes), for the first reply to arrive (to allow time for a server application to be triggered), and 15 seconds for subsequent replies, but both samples can end without getting any replies. See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of the Request sample programs.

Running the Request sample programs:

Running the amqsreq0.c, amqsreq, and amqsreqc samples

The C version of the program takes three parameters:

1. The name of the target server queue (necessary)
2. The name of the queue manager (optional)
3. The reply queue (optional)

For example, enter one of the following:

- `amqsreq myqueue qmanagername replyqueue`
- `amqsreqc myqueue qmanagername`
- `amq0req0 myqueue`

where `myqueue` is the name of the target server queue, `qmanagername` is the name of the queue manager that owns `myqueue`, and `replyqueue` is the name of the reply queue.

If you omit the name of the queue manager, it is assumed that the default queue manager owns the queue. If you omit the name of the reply queue, the default reply queue is provided.

Running the amq0req0.cbl sample

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

Please enter the name of the target server queue

The program takes its input from StdIn and adds each line to the target server queue, taking each line of text as the content of a request message. The program ends when a null line is read.

Running the AMQSREQ4 sample

The C program creates messages by taking data from stdin (the keyboard) with a blank time terminating input. The program takes up to three parameters: the name of the target queue (required), the queue manager name (optional), and the reply-to queue name (optional). If no queue manager name is specified, the default queue manager is used. If no reply-to queue is specified, the `SYSTEM.SAMPLE.REPLY` queue is used.

Here is an example of how to call the C sample program, specifying the reply-to queue, but letting the queue manager default:

```
CALL PGM(QMQM/AMQSREQ4) PARM('SYSTEM.SAMPLE.LOCAL' ' ' 'SYSTEM.SAMPLE.REPLY')
```


Note: Remember that queue names are case sensitive. All the queues created by the sample file create program `AMQSAMP4` have names created in uppercase characters.

Running the AMQ0REQ4 sample

The COBOL program creates messages by accepting data from the keyboard. To start the program, call the program and specify the name of your target queue as a parameter. The program accepts input from the keyboard into a buffer and creates a request message for each line of text. The program stops when you enter a blank line at the keyboard.

Running the Request sample using triggering:

If the sample is used with triggering and one of the Inquire, Set, or Echo sample programs, the line of input must be the queue name of the queue that you want the triggered program to access.

Running the Request sample using triggering on UNIX, Linux, and Windows: 

On UNIX, Linux, and Windows, start the trigger monitor program RUNMQTRM in one session, and then start the amqsreq program in another session.

To run the samples using triggering:

1. Start the trigger monitor program RUNMQTRM in one session (the initiation queue SYSTEM.SAMPLE.TRIGGER is available for you to use).
2. Start the amqsreq program in another session.
3. Make sure that you have defined a target server queue.

The sample queues available to you to use as the target server queue for the request sample to put messages are:

- SYSTEM.SAMPLE.INQ - for the Inquire sample program
- SYSTEM.SAMPLE.SET - for the Set sample program
- SYSTEM.SAMPLE.ECHO - for the Echo sample program

These queues have a trigger type of FIRST, so if there are already messages on the queues before you run the Request sample, server applications are not triggered by the messages you send.

4. Make sure that you have defined a queue for the Inquire, Set or Echo sample program to use.

This means that the trigger monitor is ready when the request sample sends a message.

Note: The sample process definitions created using RUNMQSC and the amqscos0.tst file trigger the C samples. Change the process definitions in amqscos0.tst and use RUNMQSC with this updated file to use COBOL versions.

Figure 87 on page 515 demonstrates how to use the Request and Inquire samples together.

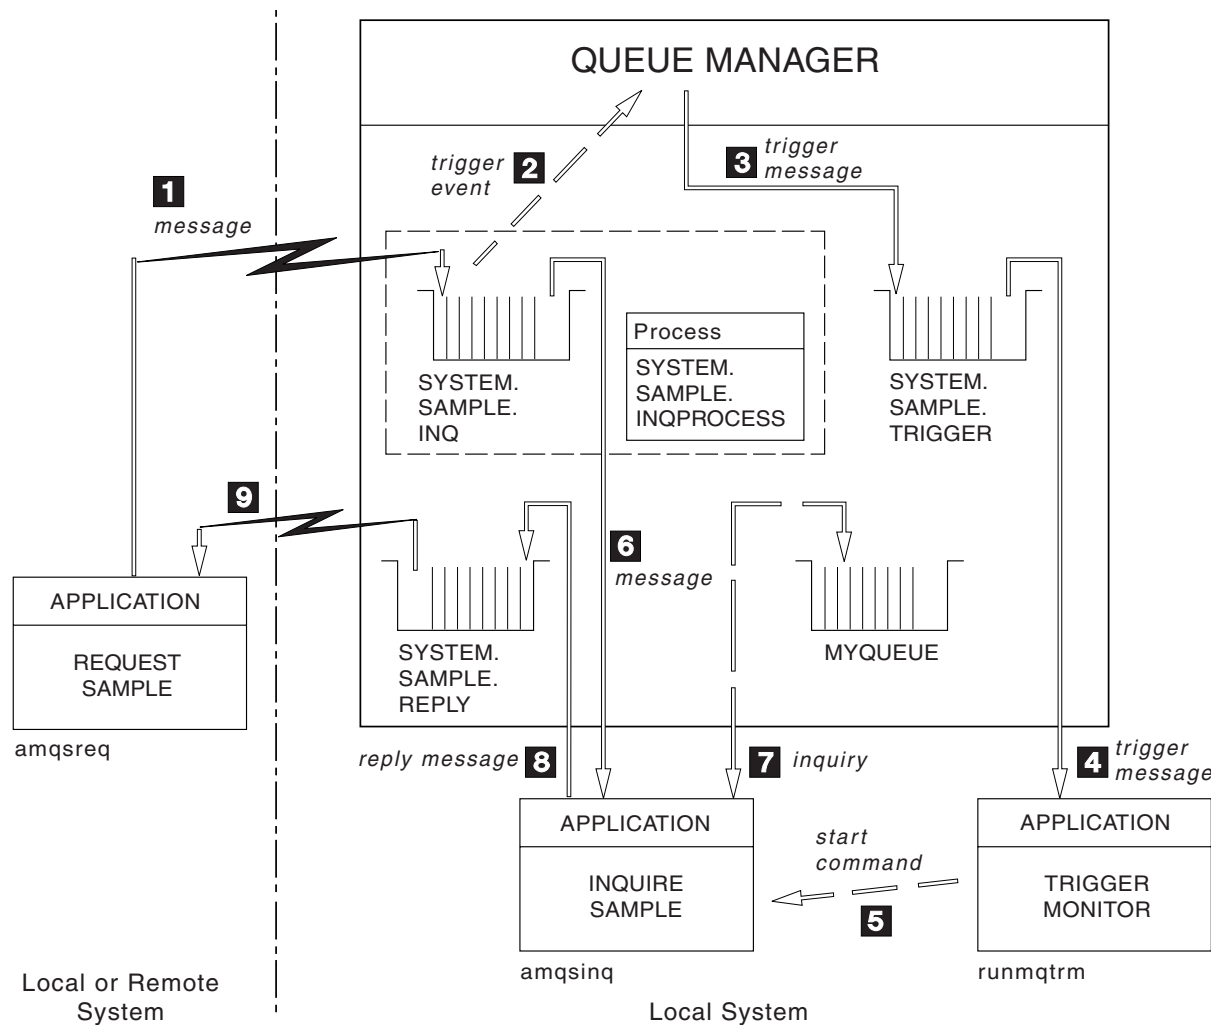


Figure 87. Request and Inquire samples using triggering

In Figure 87 the Request sample puts messages onto the target server queue, SYSTEM.SAMPLE.INQ, and the Inquire sample queries the queue, MYQUEUE. Alternatively, you can use one of the sample queues defined when you ran amqscos0.tst, or any other queue that you have defined, for the Inquire sample.

Note: The numbers in Figure 87 show the sequence of events.

To run the Request and Inquire samples, using triggering:

1. Check that the queues that you want to use are defined. Run amqscos0.tst, to define the sample queues, and define a queue MYQUEUE.
2. Run the trigger monitor command RUNMQTRM:


```
RUNMQTRM -m qmanagername -q SYSTEM.SAMPLE.TRIGGER
```
3. Run the request sample


```
amqsreq SYSTEM.SAMPLE.INQ
```


Note: The process object defines what is to be triggered. If the client and server are not running on the same platform, any processes started by the trigger monitor must define *ApplType*, otherwise the server takes its default definitions (that is, the type of application that is normally associated with the server machine) and causes a failure.

For a list of application types, see *ApplType*.

4. Enter the name of the queue that you want the Inquire sample to use:
MYQUEUE
5. Enter a blank line (to end the Request program).
6. The request sample will then display a message, containing the data the Inquire program obtained from MYQUEUE.

You can use more than one queue; in this case, enter the names of the other queues at step 4.

For more information on triggering see “Starting IBM MQ applications using triggers” on page 234.

Running the Request sample using triggering on IBM i: 

On IBM i, start the sample trigger server, AMQSERV4, in one job, then start AMQSREQ4 in another. This means that the trigger server is ready when the Request sample program sends a message.

Note:

1. The sample definitions created by AMQSAMP4 trigger the C versions of the samples. If you want to trigger the COBOL versions, change the process definitions SYSTEM.SAMPLE.ECHOPROCESS, SYSTEM.SAMPLE.INQPROCESS, and SYSTEM.SAMPLE.SETPROCESS. You can use the CHGMQMPCRC command (for details, see Change MQ Process (CHGMQMPCRC)) to do this, or edit and run your own version of AMQSAMP4.
2. Source code for AMQSERV4 is supplied for the C language only. However, a compiled version (that you can use with the COBOL samples) is supplied in library QMQM.

You could put your request messages on these sample server queues:

- SYSTEM.SAMPLE.ECHO (for the Echo sample programs)
- SYSTEM.SAMPLE.INQ (for the Inquire sample programs)
- SYSTEM.SAMPLE.SET (for the Set sample programs)

A flow chart for the SYSTEM.SAMPLE.ECHO program is shown in Figure 88 on page 518. Using the example data file the command to issue the C program request to this server is:

```
CALL PGM(QMQMSAMP/AMQSREQ4) PARM('QMQMSAMP/AMQSDATA(ECHO)')
```

Note: This sample queue has a trigger type of FIRST, so if there are already messages on the queue before you run the Request sample, server applications are not triggered by the messages you send.

If you want to attempt further examples, you can try the following variations:

- Use AMQSTRG4 (or its command line equivalent STRMQMTRM, for details, see Start MQ Trigger Monitor (STRMQMTRM)) instead of AMQSERV4 to submit the job instead, but potential job submission delays could make it less easy to follow what is happening.
- Run the SYSTEM.SAMPLE.INQUIRE and SYSTEM.SAMPLE.SET sample programs. Using the example data file the commands to issue the C program requests to these servers are, respectively:

```
CALL PGM(QMQMSAMP/AMQSREQ4) PARM('QMQMSAMP/AMQSDATA(INQ)')
CALL PGM(QMQMSAMP/AMQSREQ4) PARM('QMQMSAMP/AMQSDATA(SET)')
```

These sample queues also have a trigger type of FIRST.

Design of the Request sample program:

The program opens the target server queue so that it can put messages. It uses the MQOPEN call with the MQOO_OUTPUT option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

The program then opens the reply-to queue called SYSTEM.SAMPLE.REPLY so that it can get reply messages. For this, the program uses the MQOPEN call with the MQOO_INPUT_EXCLUSIVE option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

For each line of input, the program then reads the text into a buffer and uses the MQPUT call to create a request message containing the text of that line. On this call the program uses the MQRO_EXCEPTION_WITH_DATA report option to request that any report messages sent about the request message will include the first 100 bytes of the message data. The program continues until either it reaches the end of the input or the MQPUT call fails.

The program then uses the MQGET call to remove reply messages from the queue, and displays the data contained in the replies. The MQGET call uses the MQGMO_WAIT, MQGMO_CONVERT, and MQGMO_ACCEPT_TRUNCATED options. The *WaitInterval* is 5 minutes in the COBOL version, and 1 minute in the C version, for the first reply (to allow time for a server application to be triggered), and 15 seconds for subsequent replies. The program waits for these periods if there is no message on the queue. If no message arrives before this interval expires, the call fails and returns the MQRC_NO_MSG_AVAILABLE reason code. The call also uses the MQGMO_ACCEPT_TRUNCATED_MSG option, so that messages longer than the declared buffer size are truncated.

The program demonstrates how to clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The program continues until either the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code or the MQGET call fails. If the call fails, the program displays an error message that contains the reason code.

The program then closes both the target server queue and the reply-to queue using the MQCLOSE call.

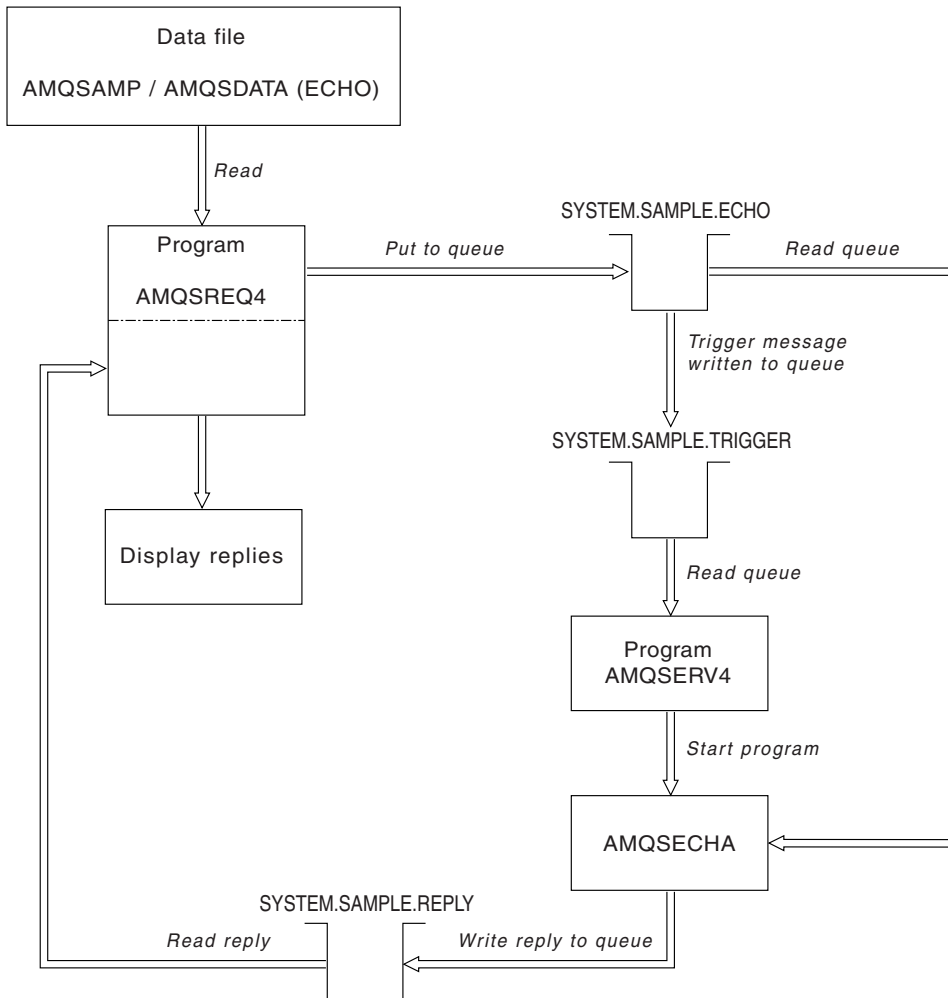


Figure 88. Sample IBM i Client/Server (Echo) program flowchart

The Set sample programs:



The Set sample programs inhibit put operations on a queue by using the MQSET call to change the queue's **InhibitPut** attribute. Also, learn about the design of Set sample programs.

See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

The programs are intended to run as triggered programs, so their only input is an MQTMC2 (trigger message) structure that contains the name of a target queue with attributes that are to be inquired upon. The C version also uses the queue manager name. The COBOL version uses the default queue manager.

For the triggering process to work, ensure that the Set sample program that you want to use is triggered by messages arriving on queue SYSTEM.SAMPLE.SET. To do this, specify the name of the Set sample program that you want to use in the *ApplicId* field of the process definition SYSTEM.SAMPLE.SETPROCESS. The sample queue has a trigger type of FIRST; if there are already messages on the queue before you run the Request sample, the Set sample is not triggered by the messages that you send.

When you have set the definition correctly:

-  For UNIX, Linux, and Windows systems, start the **runmqtrm** program in one session, then start the **amqsreq** program in another.
-  For IBM i, start the **AMQSERV4** program in one session, then start the **AMQSREQ4** program in another. You could use **AMQSTRG4** instead of **AMQSERV4**, but potential job submission delays could make it less easy to follow what is happening.

Use the Request sample programs to send request messages, each containing just a queue name, to queue **SYSTEM.SAMPLE.SET**. For each request message, the Set sample programs send a reply message containing a confirmation that put operations have been inhibited on the specified queue. The replies are sent to the reply-to queue specified in the request message.

Design of the Set sample program

The program opens the queue named in the trigger message structure that it was passed when it started. (For clarity, we will call this the *request queue*.) The program uses the **MQOPEN** call to open this queue for shared input.

The program uses the **MQGET** call to remove messages from this queue. This call uses the **MQGMO_ACCEPT_TRUNCATED_MSG** and **MQGMO_WAIT** options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message, and displays a warning message.

For each request message removed from the request queue, the program reads the name of the queue (which we will call the *target queue*) contained in the data and opens that queue using the **MQOPEN** call with the **MQOO_SET** option. The program then uses the **MQSET** call to set the value of the **InhibitPut** attribute of the target queue to **MQQA_PUT_INHIBITED**.

If the **MQSET** call is successful, the program uses the **MQPUT1** call to put a reply message on the reply-to queue. This message contains the string **PUT inhibited**.

If the **MQOPEN** or **MQSET** call is unsuccessful, the program uses the **MQPUT1** call to put a report message on the reply-to queue. In the *Feedback* field of the message descriptor of this report message is the reason code returned by either the **MQOPEN** or **MQSET** call, depending on which one failed.

After the **MQSET** call, the program closes the target queue using the **MQCLOSE** call.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

The TLS sample program:

AMQSSSLC is a sample C program that demonstrates how to use the MQCNO and MQSCO structures to supply TLS client connection information on the MQCONNX call. This enables a client MQI application to provide the definition of its client connection channel and TLS settings at run time without a client channel definition table (CCDT).

If a connection name is supplied, the program constructs a client connection channel definition in an MQCD structure.

If the stem name of the key repository file is supplied, the program constructs an MQSCO structure; if an OCSP responder URL is also supplied, the program constructs an authentication information record MQAIR structure.

The program then connects to the queue manager using MQCONNX. It inquires and prints out the name of the queue manager to which it connected.

This program is intended to be linked as an MQI client application. However, it can be linked as a regular MQI application. Then, it simply connects to a local queue manager and ignores the client connection information

AMQSSSLC accepts the following parameters, all of which are optional:

- m QmgrName**
Name of the queue manager to connect to
- c ChannelName**
Name of the channel to use
- x ConnName**
Server connection name

TLS parameters:

- k KeyReposStem**
The stem name of the key repository file. This is the full path to the file without the .kdb suffix.
For example:
/home/user/client
C:\User\client
- s CipherSpec**
The TLS channel CipherSpec string corresponding to the SSLCIPH on the SVRCONN channel definition on the queue manager.
- f** Specifies that only FIPS 140-2 certified algorithms must be used.
- b VALUE1[,VALUE2...]**
Specifies that only Suite B compliant algorithms must be used. This parameter is a comma-separated list of one or more of the following values: NONE,128_BIT,192_BIT. These values have the same meaning as those for the MQSUIITEB environment variable, and the equivalent EncryptionPolicySuiteB setting in the client configuration file SSL stanza.
- p Policy**
Specifies the certificate validation policy to be used. This can be one of the following values:
 - ANY** Apply each of the certificate validation policies supported by the secure sockets library and accept the certificate chain if any of the policies considers the certificate chain valid. This setting can be used for maximum backwards compatibility with older digital certificates which do not comply with the modern certificate standards.

RFC5280

Apply only the RFC 5280 compliant certificate validation policy. This setting provides stricter validation than the ANY setting, but rejects some older digital certificates.

The default value is ANY.

OCSP certificate revocation parameter:

-o URL

The OCSP Responder URL

Running the TLS sample program:

To run the TLS sample program you must first set up your TLS environment. You then run the sample from the command line, supplying a number of parameters.

About this task

The following instructions run the sample program using personal certificates. By varying the command you can, for example, use CA certificates and check their status using an OCSP responder. See the instructions within the sample.

Procedure

1. Create a queue manager with the name QM1. For more information, see *crtmqm*.
2. Create a key repository for the queue manager. For more information, see *Setting up a key repository on UNIX, Linux, and Windows*.
3. Create a key repository for the client. Call it *clientkey.kdb*.
4. Create a personal certificate for the queue manager. For more information, see *Creating a self-signed personal certificate on UNIX, Linux, and Windows*.
5. Create a personal certificate for the client.
6. Extract the personal certificate from the server key repository and add it to the client repository. For more information, see *Extracting the public part of a self-signed certificate from a key repository on UNIX, Linux, and Windows, and Adding a CA certificate (or the public part of a self-signed certificate) into a key repository, on UNIX, Linux or Windows systems*.
7. Extract the personal certificate from the client key repository and add it to the server key repository.
8. Create a server connection channel using the MQSC command:

```
DEFINE CHANNEL(QM1SVRCONN) CHLTYPE(SVRCONN) TRPTYPE(TCP) SSLCIPH(TLS_RSA_WITH_AES_128_CBC_SHA)
```

For more information see *Server-connection channel*

9. Define and start a channel listener on the queue manager. For more information see *DEFINE LISTENER* and *START LISTENER*.
10. Run the sample program using the following command:

```
AMQSSSLC -m QM1 -c QM1SVRCONN -x localhost  
-k "C:\Program Files\IBM\MQ\clientkey" -s TLS_RSA_WITH_AES_128_CBC_SHA  
-o http://dummy.OCSP.responder
```

Results

The sample program performs the following actions:

1. Connects to any specified queue manager, or to the default queue manager, using any options specified.
2. Opens the queue manager and inquires on its name.
3. Closes the queue manager.

4. Disconnects from the queue manager.

If the sample program runs successfully, it displays output similar to the following example:

```
Sample AMQSSSLC start
Connecting to queue manager QM1
Using the server connection channel QM1SVRCONN
on connection name localhost.
Using TLS CipherSpec TLS_RSA_WITH_AES_128_CBC_SHA
Using TLS key repository stem C:\Program Files\IBM\MQ\clientkey
Using OCSP responder URL http://dummy.OCSP.responder
Connection established to queue manager QM1
```

```
Sample AMQSSSLC end
```

If the sample program encounters a problem, it displays an appropriate error message, for example if you specify an invalid OCSP responder URL, you receive the following message:

```
MQCONN ended with reason code 2553
```

For a list of reason codes see API reason codes.

The Triggering sample programs:

The function provided in the triggering sample is a subset of that provided in the trigger monitor in the **runmqtrm** program.

See “Features demonstrated in the sample programs on Multiplatforms” on page 459 for the names of these programs.

Design of the triggering sample

The triggering sample program opens the initiation queue using the MQOPEN call with the MQOO_INPUT_AS_Q_DEF option. It gets messages from the initiation queue using the MQGET call with the MQGMO_ACCEPT_TRUNCATED_MSG and MQGMO_WAIT options, specifying an unlimited wait interval. The program clears the *MsgId* and *CorrelId* fields before each MQGET call to get messages in sequence.

When it has retrieved a message from the initiation queue, the program tests the message by checking the size of the message to make sure that it is the same size as an MQTM structure. If this test fails, the program displays a warning.

For valid trigger messages, the triggering sample copies data from these fields: *ApplicId*, *EnvrData*, *Version*, and *ApplType*. The last two of these fields are numeric, so the program creates character replacements to use in an MQTMC2 structure for IBM i, UNIX, Linux, and Windows systems.

The triggering sample issues a start command to the application specified in the *ApplicId* field of the trigger message, and passes an MQTMC2 or MQTMC (a character version of the trigger message) structure.

- ▶ **ULW** In UNIX, Linux, and Windows systems, the *EnvrData* field is used as an extension to the invoking command string.
- ▶ **IBM i** In IBM i, it is used as job submission parameters, for example, the job priority or the job description.

Finally, the program closes the initiation queue.

Ending the triggering sample programs on IBM i

A trigger monitor program can be ended by the sysrequest option 2 (ENDRQS) or by inhibiting gets from the trigger queue.

If the sample trigger queue is used, the command is:

```
CHGMQM QNAME('SYSTEM.SAMPLE.TRIGGER') MQMNAME GETENBL(*NO)
```

Important: Before starting triggering again on this queue, you must enter the following command:

```
CHGMQM QNAME('SYSTEM.SAMPLE.TRIGGER') GETENBL(*YES)
```

Running the Triggering sample programs:

This topic contains information about running Triggering sample programs.

Running the amqstrg0.c, amqstrg, and amqstrgc samples

The program takes 2 parameters:

1. The name of the initiation queue (necessary)
2. The name of the queue manager (optional)

If a queue manager is not specified, it connects to the default one. A sample initiation queue will have been defined when you ran amqscos0.tst; the name of that queue is SYSTEM.SAMPLE.TRIGGER, and you can use it when you run this program.

Note: The function in this sample is a subset of the full triggering function that is supplied in the runmqtrm program.

Running the AMQSTRG4 sample

This is a trigger monitor for the IBM i environment. It submits one IBM i job for each application to be started. This means that there is additional processing associated with each trigger message.

AMQSTRG4 (in QCSRC) takes two parameters: the name of the initiation queue that it is to serve, and the name of the queue manager (optional). AMQSAMP4 (in QCLSRC) defines a sample initiation queue, SYSTEM.SAMPLE.TRIGGER, that you can use when you try the sample programs.

Using the example trigger queue, the command to issue is:

```
CALL PGM(QMQM/AMQSTRG4) PARM('SYSTEM.SAMPLE.TRIGGER')
```

Alternatively, you can use the CL equivalent STRMQMTRM; for details, see Start MQ Trigger Monitor (STRMQMTRM).

Running the AMQSERV4 sample

This is a trigger server for the IBM i environment. For each trigger message, this server runs the start command in its own job to start the specified application. The trigger server can call CICS transactions.

AMQSERV4 takes two parameters: the name of the initiation queue that it is to serve, and the name of the queue manager (optional). AMQSAMP4 defines a sample initiation queue, SYSTEM.SAMPLE.TRIGGER, that you can use when you try the sample programs.


Using the example trigger queue the command to issue is:

```
CALL PGM(QMQM/AMQSERV4) PARM('SYSTEM.SAMPLE.TRIGGER')
```



Design of the trigger server:

The design of the trigger server is similar to that of the trigger monitor, with a few exceptions

The design of the trigger server is similar to that of the trigger monitor, except that the trigger server:

- Allows MQAT_CICS as well as MQAT_OS400 applications.
-  Calls IBM i applications in its own job (or uses STRCICSUSR to start CICS applications) rather than submitting an IBM i job.
- For CICS applications, substitutes the *EnvData*, for example, to specify the CICS region, from the trigger message in the STRCICSUSR command.
- Opens the initiation queue for shared input, so that many trigger servers can run at the same time.

Note: Programs started by AMQSERV4 must not use the MQDISC call because this stops the trigger server. If programs started by AMQSERV4 use the MQCONN call, they get the MQRC_ALREADY_CONNECTED reason code.

Using the TUXEDO samples on UNIX and Windows:  



Learn about the Put and Get sample programs for TUXEDO, and building the server environment in TUXEDO.

Before you begin

Before running these samples, you must build the server environment.

About this task

Note: Throughout this section, the backslash (\) character is used to split long commands over more than one line. Do not enter this character. Enter each command as a single line.

Building the server environment:  

Information about building the server environment for IBM MQ for different platforms.

Before you begin

It is assumed that you have a working TUXEDO environment.

How to build the server environment for IBM MQ for AIX (32-bit).

Procedure

1. Create a directory (for example, APPDIR) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO, and *MQ_INSTALLATION_PATH* represents the high-level directory in which IBM MQ is installed:

```
$ export CFLAGS="-I MQ_INSTALLATION_PATH/inc -I /APPDIR -L MQ_INSTALLATION_PATH/lib"
$ export LDOPTS="-lmqm"
$ export FIELDTBLS= MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ export VIEWFILES=/APPDIR/amqstxvx.V
$ export LIBPATH=$TUXDIR/lib: MQ_INSTALLATION_PATH/lib:/lib
```

3. Add the following line to the TUXEDO file udataobj/RM:

```
MQSeries_XA_RMI:MQRMIXASwitchDynamic: -lmqmx -lmqm
```

4. Run the commands:

```
$ mklhdhr MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ viewc MQ_INSTALLATION_PATH/samp/amqstxvx.v
$ buildtms -o MQXA -r MQSeries_XA_RMI
$ buildserver -o QSERV1 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
  -f MQ_INSTALLATION_PATH/lib/libmqm.a \
  -r MQSeries_XA_RMI -s MPUT1:MPUT \
  -s MGET1:MGET \
  -v -bshm
$ buildserver -o QSERV2 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
  -f MQ_INSTALLATION_PATH/lib/libmqm.a \
  -r MQSeries_XA_RMI -s MPUT2:MPUT \
  -s MGET2:MGET \
  -v -bshm
$ buildclient -o doputs -f MQ_INSTALLATION_PATH/samp/amqstxpx.c \
  -f MQ_INSTALLATION_PATH/lib/libmqm.a
$ buildclient -o dogets -f MQ_INSTALLATION_PATH/samp/amqstxgx.c \
  -f MQ_INSTALLATION_PATH/lib/libmqm.a
```

5. Edit ubbstxcx.cfg and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmlodcf -y MQ_INSTALLATION_PATH/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> crd1 -z /APPDIR/TLOG1
```

7. Start the queue manager:

```
$ strmqm
```

8. Start Tuxedo:

```
$ tmbboot -y
```

What to do next

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

How to build the server environment for IBM MQ for AIX (64-bit).

Procedure

1. Create a directory (for example, APPDIR) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where TUXDIR represents the root directory for TUXEDO, and `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.:

```
$ export CFLAGS="-I MQ_INSTALLATION_PATH/inc -I APPDIR -L MQ_INSTALLATION_PATH/lib64"  
$ export LDOPTS="-lmqm"  
$ export FIELDTBLS= MQ_INSTALLATION_PATH/samp/amqstxvx.flDs  
$ export VIEWFILES=APPDIR/amqstxvx.V  
$ export LIBPATH=TUXDIR/lib64: MQ_INSTALLATION_PATH/lib64:/lib64
```

3. Add the following line to the TUXEDO file `udataobj/RM`:

```
MQSeries_XA_RMI:MQRMIXASwitchDynamic: -lmqmx64 -lmqm
```

4. Run the commands:

```
$ mkfldhdr MQ_INSTALLATION_PATH/samp/amqstxvx.flDs  
$ viewc MQ_INSTALLATION_PATH/samp/amqstxvx.v  
$ buildtms -o MQXA -r MQSeries_XA_RMI  
$ buildserver -o MQSERV1 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \  
-f MQ_INSTALLATION_PATH/lib64/libmqm.a \  
-r MQSeries_XA_RMI -s MPUT1:MPUT \  
-s MGET1:MGET \  
-v -bshm  
$ buildserver -o MQSERV2 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \  
-f MQ_INSTALLATION_PATH/lib64/libmqm.a \  
-r MQSeries_XA_RMI -s MPUT2:MPUT  
-s MGET2:MGET \  
-v -bshm  
$ buildclient -o doputs -f MQ_INSTALLATION_PATH/samp/amqstxpx.c \  
-f MQ_INSTALLATION_PATH/lib64/libmqm.a  
$ buildclient -o dogets -f MQ_INSTALLATION_PATH/samp/amqstxgx.c \  
-f MQ_INSTALLATION_PATH/lib64/libmqm.a
```

5. Edit `ubbstxcx.cfg` and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmloadcf -y MQ_INSTALLATION_PATH/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$ tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> crd1 -z APPDIR/TLOG1
```

7. Start the queue manager:

```
$ strmqm
```

8. Start Tuxedo:

```
$ tmbboot -y
```

What to do next

You can now use the `doputs` and `dogets` programs to put messages to a queue and retrieve them from a queue.

How to build the server environment for IBM MQ for HP-UX (32-bit).

About this task

Note: The 32-bit TUXEDO server environment can only be built on the Itanium platform.

Procedure

1. Create a directory (for example, APPDIR) in which the server environment is built and execute all commands in this directory.

2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:

```
$ export CFLAGS="-Aa -D HPUX_SOURCE"
$ export FIELDTBLS= MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ export VIEWFILES=$APPDIR/amqstxvx.V
$ export TUXCONFIG=$APPDIR/tuxconfig
$ export PATH=$TUXDIR/bin:/usr/bin:/sbin: MQ_INSTALLATION_PATH/bin:$PATH
$ export SHLIB_PATH=$TUXDIR/lib: MQ_INSTALLATION_PATH/lib:/lib
$ export FLDTBLDIR=$APPDIR:$TUXDIR/udataobj
```

3. Add the following line to the TUXEDO file udataobj/RM:

```
MQSeries_XA_RMI:MQRMIASwitchDynamic: \
MQ_INSTALLATION_PATH/lib/libmqmxa.so MQ_INSTALLATION_PATH/lib/libmqm.so \
/opt/tuxedo/lib/libtux.sl
```

4. Run the commands:

```
$ mkfldhdr MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ viewc MQ_INSTALLATION_PATH/samp/amqstxvx.v
```

After running the mkfldhdr and viewc commands, the amqstxvx.h header file is created in the TUXEDO application directory. Copy this file from the TUXEDO application directory into the TUXEDO include directory, and then run the following commands.

```
$ buildtms -o MQXA -r MQSERIES_XA_RMI
$ buildserver -o MQSERV1 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-r MQSERIES_XA_RMI -s MPUT1:MPUT \
-s MGET1:MGET \
-v -bshm
$ buildserver -o MQSERV2 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-r MQSERIES_XA_RMI -s MPUT2:MPUT \
-s MGET2:MGET \
-v -bshm
$ buildclient -o doputs -f MQ_INSTALLATION_PATH/samp/amqstxpx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so
$ buildclient -o dogets -f MQ_INSTALLATION_PATH/samp/amqstxgx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so
```

5. Edit ubbstxcx.cfg and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmloadcf -y MQ_INSTALLATION_PATH/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$ tmadm -c
```

A prompt is displayed. At this prompt, enter:

```
> crdl -z /APPDIR/TLOG1
```

7. Start the queue manager:


```
$ strmqm
```

8. Start TUXEDO:

```
$ tmboot -y
```

What to do next

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

Building the server environment for HP-UX (64-bit): 

How to build the server environment for IBM MQ for HP-UX (64-bit).

About this task

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Procedure

1. Create a directory (for example, `APPDIR`) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where `TUXDIR` is the root directory for TUXEDO:

```
$ export CFLAGS="-Aa -D_HPUX_SOURCE"
$ export FIELDTBLS=MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ export VIEWFILES=$APPDIR/amqstxvx.V
$ export TUXCONFIG=$APPDIR/tuxconfig
$ export PATH=$TUXDIR/bin:/usr/bin:/sbin:MQ_INSTALLATION_PATH/bin:$PATH
$ export SHLIB_PATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib64:/lib64
$ export FLDTBLDIR=$APPDIR:$TUXDIR/udataobj
```

3. On the HP-UX IA64 (IPF) platform, add the following line to the TUXEDO file `udataobj/RM`:

```
MQSeries_XA_RMI:MQRMIASwitchDynamic: \
MQ_INSTALLATION_PATH/lib64/libmqmxa64.so MQ_INSTALLATION_PATH/lib64/libmqm.so \
/opt/tuxedo/lib/libtux.sl
```

Note: The IBM MQ libraries shipped on the HP-UX IA64 (IPF) platform have a `.so` file name extension.

4. Run the commands:

```
$ mkfldhdr MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ viewc MQ_INSTALLATION_PATH/samp/amqstxvx.v
```

After running the `mkfldhdr` and `viewc` commands, the `amqstxvx.h` header file is created in the TUXEDO application directory. Copy this file from the TUXEDO application directory into the TUXEDO include directory, and then run the following commands.

```
$ buildtms -o MQXA -r MQSERIES_XA_RMI
```

On the HP-UX IA64 (IPF) platform:

```
$ buildserver -o MQSERV1 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so \
-r MQSERIES_XA_RMI -s MPUT1:MPUT \
-s MGET1:MGET \
-v -bshm
$ buildserver -o MQSERV2 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so \
-r MQSERIES_XA_RMI -s MPUT2:MPUT \
-s MGET2:MGET \
-v -bshm
$ buildclient -o doputs -f MQ_INSTALLATION_PATH/samp/amqstxpx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so
$ buildclient -o dogets -f MQ_INSTALLATION_PATH/samp/amqstxgx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so
```

5. Edit `ubbstxcx.cfg` and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmloadcf -y MQ_INSTALLATION_PATH/samp/ubbstxcx.cfg
```


6. Create the TLOGDEVICE:

```
$tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> crdl -z /APPDIR/TLOG1
```

7. Start the queue manager:


```
$ strmqm
```

8. Start TUXEDO:

```
$ tmboot -y
```

What to do next

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

Building the server environment for Solaris (32-bit): 

How to build the server environment for IBM MQ for Solaris (32-bit).

About this task

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Procedure

1. Create a directory (for example, APPDIR) in which the server environment is built and execute all commands in this directory.

2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:

```
$ export CFLAGS="-I /APPDIR"
$ export FIELDTBLS=amqstxvx.fl ds
$ export VIEWFILES=amqstxvx.V
$ export SHLIB_PATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib:/lib
$ export LD_LIBRARY_PATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib:/lib
```

3. Add the following to the TUXEDO file udataobj/RM (RM must include `MQ_INSTALLATION_PATH/lib/libmqmcs` and `MQ_INSTALLATION_PATH/lib/libmqmzse`).

```
MQSeries_XA_RMI:MQRMIXASwitchDynamic: \
MQ_INSTALLATION_PATH/lib/libmqmxa.a MQ_INSTALLATION_PATH/lib/libmqm.so \
/opt/tuxedo/lib/libtux.a MQ_INSTALLATION_PATH/lib/libmqmcs.so \
MQ_INSTALLATION_PATH/lib/libmqmzse.so
```

4. Run the commands:

```
$ mkfldhdr amqstxvx.fl ds
$ viewc amqstxvx.v
$ buildtms -o MQXA -r MQSERIES_XA_RMI
$ buildserver -o MQSERV1 -f amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-r MQSERIES_XA_RMI -s MPUT1:MPUT \
-s MGET1:MGET \
-v -bshm
-l -ldl
$ buildserver -o MQSERV2 -f amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-r MQSERIES_XA_RMI -s MPUT2:MPUT \
-s MGET2:MGET \
-v -bshm
-l -ldl
$ buildclient -o doputs -f amqstxpx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-f MQ_INSTALLATION_PATH/lib/libmqmzse.co \
-f MQ_INSTALLATION_PATH/lib/libmqmcs.so
```

- ```
$ buildclient -o dogets -f amqstxgx.c \
 -f MQ_INSTALLATION_PATH/lib/libmqm.so
 -f MQ_INSTALLATION_PATH/lib/libmqmzse.co \
 -f MQ_INSTALLATION_PATH/lib/libmqmcs.so
```
5. Edit ubbstxcx.cfg and add details of the machine name, working directories, and queue manager as necessary:
 

```
$ tmloadcf -y ubbstxcx.cfg
```
  6. Create the TLOGDEVICE:
 

```
$ tmadmin -c
```

 A prompt is displayed. At this prompt, enter:
 


```
> crdl -z /APPDIR/TLOG1
```
  7. Start the queue manager:
 

```
$ strmqm
```
  8. Start Tuxedo:
 

```
$ tmboot -y
```

## What to do next

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

Building the server environment for Solaris (64-bit): 

How to build the server environment for IBM MQ for Solaris (64-bit).

## About this task

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

## Procedure

1. Create a directory (for example, APPDIR) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:
 

```
$ export CFLAGS="-I /APPDIR"
$ export FIELDTBLS=amqstxvx.flds
$ export VIEWFILES=amqstxvx.V
$ export SHLIB_PATH=$TUXDIR/lib: MQ_INSTALLATION_PATH/lib:/lib64
$ export LD_LIBRARY_PATH=$TUXDIR/lib64: MQ_INSTALLATION_PATH/lib64:/lib64
```
3. Add the following to the TUXEDO file udataobj/RM (RM must include `MQ_INSTALLATION_PATH/lib/libmqmcs` and `MQ_INSTALLATION_PATH/lib/libmqmzse`).
 

```
MQSERIES_XA_RMI:MQRMIXASwitchDynamic: \
MQ_INSTALLATION_PATH/lib64/libmqmxa64.a MQ_INSTALLATION_PATH/lib64/libmqm.so \
/opt/tuxedo/lib64/libtux.a MQ_INSTALLATION_PATH/lib64/libmqmcs.so \
MQ_INSTALLATION_PATH/lib64/libmqmzse.so
```
4. Run the commands:
 

```
$ mkfldhdr amqstxvx.flds
$ viewc amqstxvx.v
$ buildtms -o MQXA -r MQSERIES_XA_RMI
$ buildserver -o MQSERV1 -f amqstxsx.c \
 -f MQ_INSTALLATION_PATH/lib64/libmqm.so \
 -r MQSERIES_XA_RMI -s MPUT1:MPUT \
 -s MGET1:MGET \
 -v -bshm
 -l -ldl
$ buildserver -o MQSERV2 -f amqstxsx.c \
```

```

-f MQ_INSTALLATION_PATH/lib64/libmqm.so \
-r MQSERIES_XA_RMI -s MPUT2:MPUT \
-s MGET2:MGET \
-v -bshm
-l -ldl
$ buildclient -o doputs -f amqstpx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so \
-f MQ_INSTALLATION_PATH/lib64/libmqmzse.co \
-f MQ_INSTALLATION_PATH/lib64/libmqmcs.so
$ buildclient -o dogets -f amqstxgx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so
-f MQ_INSTALLATION_PATH/lib64/libmqmzse.co \
-f MQ_INSTALLATION_PATH/lib64/libmqmcs.so

```

5. Edit `ubbstxcx.cfg` and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmloadcf -y ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$ tadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> crdl -z /APPPDIR/TLOG1
```

7. Start the queue manager:

```
$ strmqm
```

8. Start Tuxedo:

```
$ tmbot -y
```

## What to do next

You can now use the `doputs` and `dogets` programs to put messages to a queue and retrieve them from a queue.

Building the server environment for Windows (32-bit): **Windows**

Building the server environment for IBM MQ for Windows (32-bit).

## About this task

**Note:** Change the fields identified as *VARIABLES* in the following, to the directory paths:

Table 77. Fields to change to directory paths

| Field          | Directory path                                                                                             |
|----------------|------------------------------------------------------------------------------------------------------------|
| <i>MQMDIR</i>  | The directory path specified when IBM MQ was installed, for example <code>g:\Program Files\IBM\MQ</code> . |
| <i>TUXDIR</i>  | The directory path specified when TUXEDO was installed, for example <code>f:\tuxedo</code> .               |
| <i>APPPDIR</i> | The directory path to be used for the sample application, for example <code>f:\tuxedo\apps\mqapp</code> .  |

```

*RESOURCES
IPCKEY 99999
UID 0
GID 0
MAXACCESSERS 20
MAXSERVERS 20
MAXSERVICES 50
MASTER SITE1
MODEL SHM
LDBAL N

*MACHINES
MachineName LMID=SITE1
 TUXDIR="f:\tuxedo"
 APPDIR="f:\tuxedo\apps\mqapp;g:\Program Files\IBM\WebSphere MQ\bin"
 ENVFILE="f:\tuxedo\apps\mqapp\amqstxen.env"
 TUXCONFIG="f:\tuxedo\apps\mqapp\tuxconfig"
 ULOGPFX="f:\tuxedo\apps\mqapp\ULOG"
 TLOGDEVICE="f:\tuxedo\apps\mqapp\TLOG"
 TLOGNAME=TLOG
 TYPE="i386NT"
 UID=0
 GID=0

*GROUPS
GROUP1
 LMID=SITE1 GRPNO=1
 TMSNAME=MQXA
 OPENINFO="MQSERIES_XA_RMI:MYQUEUEMANAGER"

*SERVERS
DEFAULT: CLOPT="-A -- -m MYQUEUEMANAGER"

MQSERV1 SRVGRP=GROUP1 SRVID=1
MQSERV2 SRVGRP=GROUP1 SRVID=2

*SERVICES
MPUT1
MGET1
MPUT2
MGET2

```

Figure 89. Example of *ubbstxcn.cfg* file for IBM MQ for Windows

**Note:** Change the machine name *MachineName*, and the directory paths, to match your installation. Also change the queue manager name *MYQUEUEMANAGER* to the name of the queue manager that you want to connect to.

The sample *ubbstxcn.cfg* file for IBM MQ for Windows is listed in Figure 89. It is supplied as *ubbstxcn.cfg* in the IBM MQ samples directory.

The sample makefile (see Figure 90 on page 533) supplied for IBM MQ for Windows is called *ubbstxmn.mak*, and is held in the IBM MQ samples directory.

```

TUXDIR = f:\tuxedo
MQMDIR = g:\Program Files\IBM\WebSphere MQ
APPDIR = f:\tuxedo\apps\mqapp
MQMLIB = $(MQMDIR)\tools\lib
MQMINC = $(MQMDIR)\tools\c\include
MQMSAMP = $(MQMDIR)\tools\c\samples
INC = -f "-I$(MQMINC) -I$(APPDIR)"
DBG = -f "/Zi"

amqstx.exe:
$(TUXDIR)\bin\mkfldhdr -d$(APPDIR) $(MQMSAMP)\amqstxvx.fld
$(TUXDIR)\bin\viewc -d$(APPDIR) $(MQMSAMP)\amqstxvx.v
$(TUXDIR)\bin\buildtms -o MQXA -r MQSERIES_XA_RMI
$(TUXDIR)\bin\buildserver -o MQSERV1 -f $(MQMSAMP)\amqstxsx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG) \
-r MQSERIES_XA_RMI \
-s MPUT1:MPUT -s MGET1:MGET
$(TUXDIR)\bin\buildserver -o MQSERV2 -f $(MQMSAMP)\amqstxsx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG) \
-r MQSERIES_XA_RMI \
-s MPUT2:MPUT -s MGET2:MGET
$(TUXDIR)\bin\buildclient -o doputs -f $(MQMSAMP)\amqstxpx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG)
$(TUXDIR)\bin\buildclient -o dogets -f $(MQMSAMP)\amqstxgx.c \
-f $(MQMLIB)\mqm.lib $(INC) -v $(DBG)
$(TUXDIR)\bin\tmloadcf -y $(APPDIR)\ubbstxcn.cfg

```

Figure 90. Sample TUXEDO makefile for IBM MQ for Windows

To build the server environment and samples, complete the following steps.

### Procedure

1. Create an application directory in which to build the sample application, for example:  
f:\tuxedo\apps\mqapp
2. Copy the following sample files from the IBM MQ sample directory to the application directory:
  - amqstxmn.mak
  - amqstxen.env
  - ubbstxcn.cfg
3. Edit each of these files to set the directory names and directory paths used on your installation.
4. Edit ubbstxcn.cfg (see Figure 89 on page 532) to add details of the machine name and the queue manager that you want to connect to.
5. Add the following line to the TUXEDO file *TUXDIR*rudataobj\rm:

```
MQSERIES_XA_RMI;MQRMIXASwitchDynamic;MQMDIR\tools\lib\mqmxa.lib MQMDIR\tools\lib\mqm.lib
```

The new entry must be one line in the file.


6. Set the following environment variables:

```
TUXDIR=TUXDIR
TUXCONFIG=APPDIR\tuxconfig
FIELDTBLs=MQMDIR\tools\c\samples\amqstxvx.fld
LANG=C
```
7. Create a TLOG device for TUXEDO. To do this, invoke `tadmin -c`, and enter the following command:

```
crdl -z APPDIR\TLOG
```
8. Set the current directory to *APPDIR*, and invoke the sample makefile `amqstxmn.mak` as an external project makefile. For example, with Microsoft Visual C++ , issue the following command:

```
msvc amqstxmn.mak
```

Select **build** to build all the sample programs.

Building the server environment for Windows (64-bit): 

How to build the server environment for IBM MQ for Windows (64-bit).

### About this task

**Note:** Change the fields identified as *VARIABLES* in the following, to the directory paths:

*Table 78. Fields to change to directory paths*

| Field         | Directory path                                                                               |
|---------------|----------------------------------------------------------------------------------------------|
| <i>MQMDIR</i> | The directory path specified when IBM MQ was installed, for example g:\Program Files\IBM\MQ. |
| <i>TUXDIR</i> | The directory path specified when TUXEDO was installed, for example f:\tuxedo.               |
| <i>APPDIR</i> | The directory path to be used for the sample application, for example f:\tuxedo\apps\mqapp.  |

```

*RESOURCES
IPCKEY 99999
UID 0
GID 0
MAXACCESSERS 20
MAXSERVERS 20
MAXSERVICES 50
MASTER SITE1
MODEL SHM
LDBAL N

*MACHINES
MachineName LMID=SITE1
 TUXDIR="f:\tuxedo"
 APPDIR="f:\tuxedo\apps\mqapp;g:\Programi½Files\IBM\WebSphere MQ\bin"
 ENVFILE="f:\tuxedo\apps\mqapp\amqstxen.env"
 TUXCONFIG="f:\tuxedo\apps\mqapp\tuxconfig"
 ULOGPFX="f:\tuxedo\apps\mqapp\ULOG"
 TLOGDEVICE="f:\tuxedo\apps\mqapp\TLOG"
 TLOGNAME=TLOG
 TYPE="i386NT"
 UID=0
 GID=0

*GROUPS
GROUP1
 LMID=SITE1 GRPNO=1
 TMSNAME=MQXA
 OPENINFO="MQSERIES_XA_RMI:MYQUEUEMANAGER"

*SERVERS
DEFAULT: CLOPT="-A -- -m MYQUEUEMANAGER"

MQSERV1 SRVGRP=GROUP1 SRVID=1
MQSERV2 SRVGRP=GROUP1 SRVID=2

*SERVICES
MPUT1
MGET1
MPUT2
MGET2

```

Figure 91. Example of *ubbstxcn.cfg* file for IBM MQ for Windows

**Note:** Change the machine name *MachineName*, and the directory paths, to match your installation. Also change the queue manager name *MYQUEUEMANAGER* to the name of the queue manager that you want to connect to.

The sample *ubbbconfig* file for IBM MQ for Windows is listed in Figure 91. It is supplied as *ubbstxcn.cfg* in the IBM MQ samples directory.

The sample makefile (see Figure 92 on page 536 ) supplied for IBM MQ for Windows is called *ubbstxmn.mak*, and is held in the IBM MQ samples directory.

```

TUXDIR = f:\tuxedo
MQMDIR = g:\Program Files\IBM\WebSphere MQ
APPDIR = f:\tuxedo\apps\mqapp
MQMLIB = $(MQMDIR)\tools\lib64
MQMINC = $(MQMDIR)\tools\c\include
MQMSAMP = $(MQMDIR)\tools\c\samples
INC = -f "-I$(MQMINC) -I$(APPDIR)"
DBG = -f "/Zi"

amqstx.exe:
$(TUXDIR)\bin\mkfldhdr -d$(APPDIR) $(MQMSAMP)\amqstxvx.fld
$(TUXDIR)\bin\viewc -d$(APPDIR) $(MQMSAMP)\amqstxvx.v
$(TUXDIR)\bin\buildtms -o MQXA -r MQSERIES_XA_RMI
$(TUXDIR)\bin\buildserver -o MQSERV1 -f $(MQMSAMP)\amqstxsx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG) \
-r MQSERIES_XA_RMI \
-s MPUT1:MPUT -s MGET1:MGET
$(TUXDIR)\bin\buildserver -o MQSERV2 -f $(MQMSAMP)\amqstxsx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG) \
-r MQSERIES_XA_RMI \
-s MPUT2:MPUT -s MGET2:MGET
$(TUXDIR)\bin\buildclient -o doputs -f $(MQMSAMP)\amqstxpx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG)
$(TUXDIR)\bin\buildclient -o dogets -f $(MQMSAMP)\amqstxgx.c \
-f $(MQMLIB)\mqm.lib $(INC) -v $(DBG)
$(TUXDIR)\bin\tmloadcf -y $(APPDIR)\ubbstxcn.cfg

```

Figure 92. Sample TUXEDO makefile for IBM MQ for Windows

To build the server environment and samples, complete the following steps.

### Procedure

1. Create an application directory in which to build the sample application, for example:  
f:\tuxedo\apps\mqapp
2. Copy the following sample files from the IBM MQ sample directory to the application directory:
  - amqstxmn.mak
  - amqstxen.env
  - ubbstxcn.cfg
3. Edit each of these files to set the directory names and directory paths used on your installation.
4. Edit ubbstxcn.cfg (see Figure 91 on page 535 ) to add details of the machine name and the queue manager that you want to connect to.
5. Add the following line to the TUXEDO file *TUXDIR*dataobj\rm  
MQSERIES\_XA\_RMI;MQRMIXASwitchDynamic;MQMDIR\tools\lib64\mqmxa64.lib MQMDIR\tools\lib64\mqm.lib

The new entry must be one line in the file.

6. Set the following environment variables:  
TUXDIR=*TUXDIR*  
TUXCONFIG=*APPDIR*\tuxconfig  
FIELDTBLs=*MQMDIR*\tools\c\samples\amqstxvx.fld  
LANG=C
7. Create a TLOG device for TUXEDO. To do this, invoke *tmadmin -c*, and enter the command:  
crdl -z *APPDIR*\TLOG
8. Set the current directory to *APPDIR*, and invoke the sample makefile *amqstxmn.mak* as an external project makefile. For example, with Microsoft Visual C++ , issue the following command:  
msvc amqstxmn.mak

Select **build** to build all the sample programs.



Sample server program for TUXEDO: **UNIX** **Windows**

The sample server program (amqstxsx) is designed to run with the Put (amqstxpx.c) and the Get (amqstxgx.c) sample programs. The sample server program runs automatically when TUXEDO is started.

**Note:** You must start your queue manager before you start TUXEDO.

The sample server provides two TUXEDO services, MPUT1 and MGET1:

- The MPUT1 service is driven by the PUT sample and uses MQPUT1 in syncpoint to put a message in a unit of work controlled by TUXEDO. It takes the parameters QName and Message Text, which are supplied by the PUT sample.
- The MGET1 service opens and closes the queue each time that it gets a message. It takes the parameters QName and Message Text, which are supplied by the GET sample.

Any error messages, reason codes, and status messages are written to the TUXEDO log file.

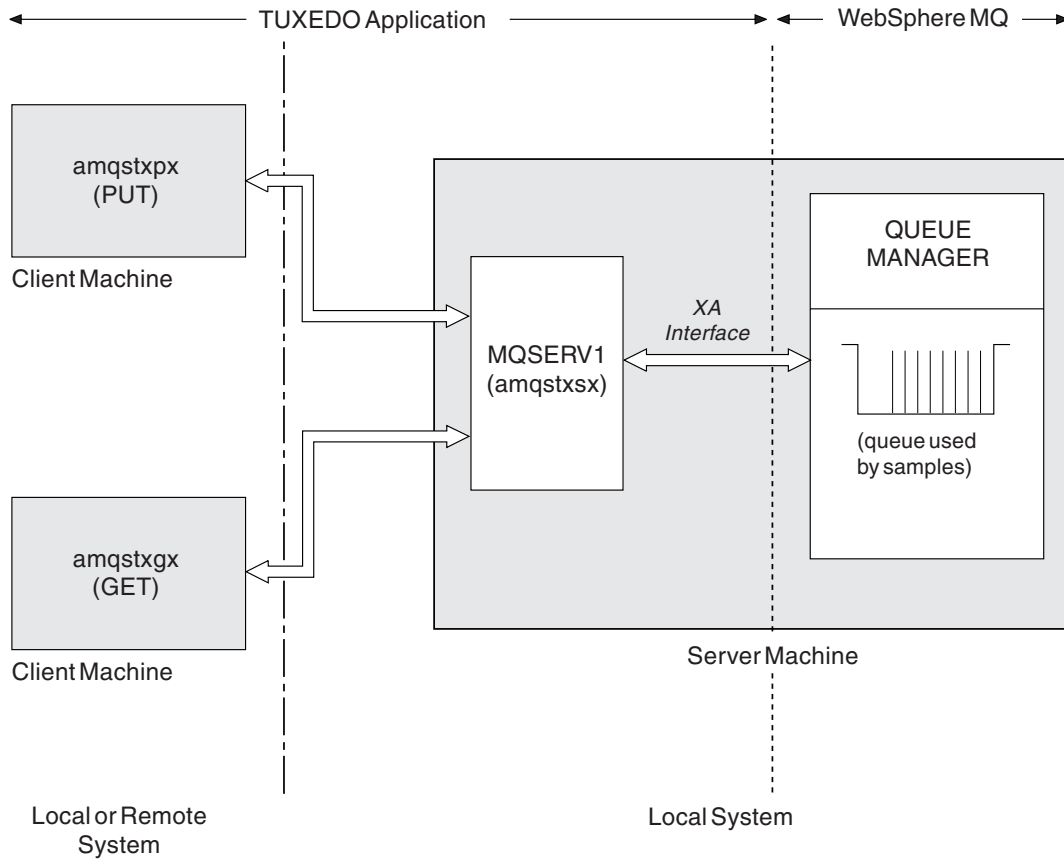


Figure 93. How TUXEDO samples work together

Put sample program for TUXEDO:  

This sample allows you to put a message on a queue multiple times, in batches, demonstrating syncpointing using TUXEDO as the resource manager.

The sample server program `amqstxsx` must be running for the put sample to succeed; the server sample program connects to the queue manager and uses the XA interface. To run the sample enter:

- `doputs -n queueName -b batchSize -c tranCount -t message`

For example:

- `doputs -n myqueue -b 5 -c 6 -t "Hello World"`

This puts 30 messages onto the queue named `myqueue`, in six batches, each with five messages in it. If there are any problems it backs a batch of messages out, otherwise it commits them.

Any error messages are written to the TUXEDO log file and to `stderr`. Any reason codes are written to `stderr`.

Get sample for TUXEDO:  

This sample allows you to get messages from a queue in batches.

The sample server program `amqstxsx` must be running for the Get sample to succeed; the sample server program connects to the queue manager and uses the XA interface. To run the sample, enter the following command:


- `dogets -n queueName -b batchSize -c tranCount`

For example:

- `dogets -n myqueue -b 6 -c 4`

This takes 24 messages off the queue named `myqueue`, in six batches, each with four messages in it. If you run this after the put example, which puts 30 messages on `myqueue`, you have only six messages on `myqueue`. The number of batches and the batch size can vary between putting the messages and getting them.

Any error messages are written to the TUXEDO log file and to `stderr`. Any reason codes are written to `stderr`.

Using the SSPI security exit on Windows: 

This topic describes how to use the SSPI channel-exit programs on Windows systems. The supplied exit code is in two formats: object and source.

### Object code

The object code file is called `amqrspin.dll`. For both client and server, it is installed as a standard part of IBM MQ for Windows in the `MQ_INSTALLATION_PATH/exits/INSTALLATION_NAME` folder. For example, `C:\Program Files\IBM\MQ\exits\installation2`. It is loaded as a standard user exit. You can run the supplied security channel exit and use authentication services in your definition of the channel.

To do this, specify either of the following:

```
SCYEXIT('amqrspin(SCY_KERBEROS)')
```

```
SCYEXIT('amqrspin(SCY_NTLM)')
```

To provide support for a restricted channel, specify the following on the SRVCONN channel:

```
SCYDATA('remote_principal_name')
```

where *remote\_principal\_name* is in the form *DOMAIN\user*. The secure channel is established only if the name of the remote principal matches *remote\_principal\_name*.

To use the supplied channel-exit programs between systems that operate within a Kerberos security domain, create a **servicePrincipalName** for the queue manager.

### Source code

The exit source code file is called `amqsspin.c`. It is in `C:\Program Files\IBM\MQ\Tools\c\Samples`.

If you modify the source code, you must recompile the modified source.

You compile and link it in the same way as any other channel exit for the relevant platform, except that SSPI headers need to be accessed at compile time, and the SSPI security libraries, together with any recommended associated libraries, need to be accessed at link time.

Before you execute the following command, make sure that `cl.exe`, and the Visual C++ library, and the include folder are available in your path. For example:

```
cl /VERBOSE /LD /MT /Ipath_to_Microsoft_platform_SDK\include
/Ipath_to_IBM_MQ\tools\c\include amqsspin.c /DSECURITY_WIN32
-link /DLL /EXPORT:SCY_KERBEROS /EXPORT:SCY_NTLM STACK:8192
```

**Note:** The source code does not include any provision for tracing or error handling. If you modify and use the source code, add your own tracing and error-handling routines.

### Running the samples using remote queues:

You can demonstrate remote queuing by running the samples on connected queue managers.

Program `amqscos0.tst` provides a local definition of a remote queue (`SYSTEM.SAMPLE.REMOTE`) that uses a remote queue manager named `OTHER`. To use this sample definition, change `OTHER` to the name of the second queue manager that you want to use. You must also set up a message channel between your two queue managers; for information on how to do this, see *Defining the channels*.

The Request sample programs put their own local queue manager name in the *ReplyToQMgr* field of messages that they send. The Inquire and Set samples send reply messages to the queue and message queue manager named in the *ReplyToQ* and *ReplyToQMgr* fields of the request messages that they process.

## The Cluster Queue Monitoring sample program (AMQSCLM):

This sample uses the built-in IBM MQ cluster workload balancing features to direct messages to instances of queues that have consuming applications attached. This automatic direction prevents the build-up of messages on an instance of a cluster queue to which no consuming application is attached.

### Overview

You can set up a cluster that has more than one definition for the same queue on different queue managers. This configuration provides the benefit of increased availability and workload balancing. However, there is no capability built into IBM MQ to dynamically modify the distribution of messages across a cluster based on the state of attached applications. For this reason, a consuming application must always be attached to every instance of a queue to ensure that messages are processed.

The cluster queue monitoring sample program monitors the state of attached applications. The program dynamically adjusts the built-in workload balancing configuration to direct messages to instances of a clustered queue with consuming applications attached. In certain situations this program can be used to relax the need for a consuming application to always be connected to every instance of a queue. It also resends messages that become queued on an instance of a queue with no consuming applications attached. Resending messages enables messages to be routed around a consuming application that is temporarily shut down.

The program is designed to be used where the consuming applications are long running applications, rather than frequently attaching and detaching applications.

The cluster queue monitoring sample program is the compiled executable program of the C sample file `amqsc1ma.c`.

Further information about clusters and workload can be found in *Using clusters for workload management*

*AMQSCLM: Design and Planning for using the sample:*

Information about how the cluster queue monitoring sample program works, points to consider when setting up a system for the sample program to run on, and modifications that can be made to the sample source code.

### Design

The cluster queue monitoring sample program monitors local clustered queues that have consuming applications attached. The program monitors queues specified by the user. The name of the queue might be specific, for example `APP.TEST01`, or generic. Generic names must be in a format that conforms to PCF (Programmable Command Format). Examples of generic names are `APP.TEST*`, or `APP*`.

Each queue manager in a cluster that owns an instance of a local queue to be monitored, requires an instance of the cluster queue monitoring sample program to be connected to it.

### Dynamic message routing

The cluster queue monitoring sample program uses the **IPPROCS** (open for input process count) value of a queue to determine whether that queue has any consumers. A value greater than 0 indicates that the queue has at least one consuming application attached. Such queues are active. A value of 0 indicates that the queue has no attached consuming programs. Such queues are inactive.

For a clustered queue with multiple instances in a cluster, IBM MQ uses the cluster workload priority property **CLWLPRTY** of each queue instance to determine which instances to send messages to. IBM MQ sends messages to the available instances of a queue with the highest **CLWLPRTY** value.

The cluster queue monitoring sample program activates a cluster queue by setting the local **CLWLPRTY** value to 1. The program deactivates a cluster queue by setting its **CLWLPRTY** value to 0.

IBM MQ clustering technology propagates the updated **CLWLPRTY** property of a clustered queue to all relevant queue managers in the cluster. For example,

- A queue manager with a connected application that puts messages to the queue.
- A queue manager that owns a local queue of the same name in the same cluster.

The propagation is done using the full repository queue managers of the cluster. New messages for the cluster queue are directed to the instances with the highest **CLWLPRTY** value within the cluster.

### Queued message transfer

The dynamic modification of the value of **CLWLPRTY** influences the routing of new messages. This dynamic modification does not affect messages already queued on a queue instance with no attached consumers, or messages that have been through the workload balancing mechanism before a modified **CLWLPRTY** value was propagated across the cluster. As a result, messages remain on any inactive queue and not be processed by a consuming application. To solve this, the cluster queue monitoring sample program is able to get messages from a local queue with no consumers, and send these messages to remote instances of the same queue where consumers are attached.

The cluster queue monitoring sample program transfers messages from an inactive local queue to one or more active remote queues by getting messages (using **MQGET** ) and putting messages (using **MQPUT** ) to the same clustered queue. This transfer causes the IBM MQ cluster workload management to select a different target instance, based on a higher **CLWLPRTY** value than that of the local queue instance. Message persistence and context are preserved during the message transfer. Message order, and any binding options are not preserved.

### Planning

The cluster queue monitoring sample program modifies the cluster configuration when there is a change in the connectivity of consuming applications. Modifications are transmitted from the queue managers where the cluster queue monitoring sample program is monitoring queues, to the full repository queue managers in the cluster. The full repository queue managers process the configuration updates and resend them to all relevant queue managers in the cluster. Relevant queue managers include those queue managers that own clustered queues of the same name (where an instance of the cluster queue monitoring sample program is running), and any queue manager where an application opened the cluster queue to put messages to it in the last 30 days.

Changes are asynchronously processed across the cluster. Therefore, after each change, different queue managers in the cluster might have different views of the configuration for a period of time.

The cluster queue monitoring sample program is only suitable for systems where consuming applications infrequently attach or detach; for example, long running consuming applications. When used to monitor systems where consuming applications are only attached for short periods, the latency incurred when distributing the configuration updates might result in queue managers in the cluster having an incorrect view of the queues where consumers are attached. This latency might result in incorrectly routed messages.

When monitoring many queues, a relatively low rate of change in attached consumers across all queues might increase cluster configuration traffic across the cluster. Increased cluster configuration traffic can result in excessive load on one or more of the following queue managers.

- The queue managers where the cluster queue monitoring sample program is running
- The full repository queue managers
- A queue manager with a connected application that puts messages to the queue
- A queue manager that owns a local queue of the same name in the same cluster

Processor usage on the full repository queue managers must be assessed. Additional processor usage is visible as message traffic on the full repository queue `SYSTEM.CLUSTER.COMMAND.QUEUE`. If messages build up on that queue, it indicates that the full repository queue managers are unable to keep up with the rate of cluster configuration change in the system.

When many queues are being monitored by the cluster queue monitoring sample program, there is an amount of work performed by the sample program and the queue manager. This work is performed, even when there are no changes to the attached consumers. The `-i` argument can be modified to reduce processor usage of the sample program on the local system, by decreasing the frequency of the monitoring cycle.

To help detect excessive activity, the cluster queue monitoring sample program reports average processing time per polling interval, elapsed processing time, and number of configuration changes. The reports are delivered in an information message, **CLM0045I**, every 30 minutes, or every 600 poll intervals, whichever is sooner.

### Cluster queue monitoring usage requirements

The cluster queue monitoring sample program has requirements and restrictions. You can modify the sample source code provided to change some of these restrictions in how it can be used. Examples listed in this section detail modifications that can be made.

- The cluster queue monitoring sample program is designed to be used to monitor queues where consuming applications are either attached, or not attached. If the system has consuming applications that are frequently attaching and detaching, the sample program might generate excessive cluster configuration activity across the entire cluster. This might have an impact on the performance of the queue managers in the cluster.
- The cluster queue monitoring sample program depends upon the underlying IBM MQ system and cluster technology. The number of queues being monitored, the frequency of monitoring and the frequency of the change of the state of each queue affects the load on the overall system. These factors must be considered when selecting the queues to be monitored and the poll interval of the monitoring.
- An instance of the cluster queue monitoring sample program must be connected to every queue manager in the cluster that owns an instance of a queue to be monitored. It is not necessary to connect the sample program to queue managers in the cluster that do not own the queues.
- The cluster queue monitoring sample program must be run with suitable authorization to access all of the IBM MQ resources required. For example,
  - The queue manager to be connected to
  - The `SYSTEM.ADMIN.COMMAND.QUEUE`
  - All the queues to be monitored when message transfer is performed
- The command server must be running for each queue manager with the cluster queue monitoring sample program connected.
- Each instance of the cluster queue monitoring sample program requires exclusive use of a local (non-clustered) queue on the queue manager that it is connected to. This local queue is used to control the sample program, and receive reply messages from inquires made to the command server of the queue manager.
- All queues to be monitored by a single instance of the cluster queue monitoring sample program must be in the same cluster. If a queue manager has queues in multiple clusters that require monitoring, multiple instances of the sample program are required. Each instance needs a local queue for control and reply messages.

- All queues to be monitored must be in a single cluster. Queues configured to use a cluster namelist are not monitored.
- Enabling the transfer of messages from inactive queues is optional. It applies to all queues being monitored by the instance of the cluster queue monitoring sample program. If only a subset of the queues being monitored require message transfer enabled, two instances of the cluster queue monitoring sample program are needed. One sample program has message transfer enabled, and the other has message transfer disabled. Each instance of the sample program needs a local queue for control and reply messages.
- IBM MQ cluster workload balancing will, by default, send messages to instances of clustered queues that reside on the same queue manager that a putting application is connected to. This must be disabled while the local queue is inactive in the following circumstances:
  - Putting applications connect to queue managers that own instances of an inactive queue that are being monitored
  - Queued messages are being transferred from inactive queues to active queues.

The local workload balancing preference on the queue can be disabled statically, through setting the `CLWLUSEQ` value to `ANY`. In this configuration messages put on local queues are distributed to local and remote queue instances to balance workload, even when there are local consuming applications. Alternatively, the cluster queue monitoring sample program can be configured to temporarily set the `CLWLUSEQ` value to `ANY` while the queue has no attached consumers which results in only local messages going to local instances of a queue while that queue is active.

- The IBM MQ system and applications must not use `CLWLPRTY` for the queues to be monitored, or channels being used. Otherwise, the actions of the cluster queue monitoring sample program on `CLWLPRTY` queue attributes might have undesired effects.
- The cluster queue monitoring sample program logs runtime information to a set of report files. A directory to store these reports is required, and the cluster queue monitoring sample program must have authorization to write to it.

*AMQSCLM: Preparing and running the sample:*

The cluster queue monitoring sample can either be run locally connected to a queue manager, or as a client connected over a channel. The sample should be running whenever the queue manager is running, when running locally it can be configured as a queue manager service to automatically start and stop the sample with the queue manager.

## Before you begin

The following steps must be completed before running the cluster queue monitoring sample.

1. Create a working queue on each queue manager for the internal use of the sample.  
Each instance of the sample needs a local non-cluster queue for exclusive internal use. You can choose the name of the queue. The example uses the name `AMQSCLM.CONTROL.QUEUE`. For example, on Windows, you can create this queue by using the following **MQSC** command:

```
DEFINE QLOCAL(AMQSCLM.CONTROL.QUEUE)
```

You can leave the values of `MAXDEPTH` and `MAXMSGL` as default.

2. Create a directory for error and information message logs.  
The sample writes diagnostic messages to report files. You must choose a directory in which to store the files. For example, on Windows, you can create a directory using the following command:

```
mkdir C:\AMQSCLM\rpts
```

The report files created by the sample have the following naming convention:

```
QmgrName.ClusterName.RPT0n.LOG
```

3. (Optional) Define the cluster queue monitoring sample as an IBM MQ service.

To monitor queues, the sample must always be running. To ensure that the cluster queue monitoring sample is always running, you can define the sample as a queue manager service. Defining the sample as a service means that AMQSCLM is started when the queue manager starts. You can use the following example to define the cluster queue monitoring sample as an IBM MQ service.

```
define service(AMQSCLM) +
 descr('Active Cluster Queue Message Distribution Monitor - AMQSCLM') +
 control(qmgr) +
 servtype(server) +
 startcmd('MQ_INSTALLATION_PATH\tools\c\samples\Bin\AMQSCLM.exe') +
 startarg('-m +QMNAME+ -c CLUSTER1 -q ABC* -r AMQSCLM.CONTROL.QUEUE -l c:\AMQSCLM\rpts') +
 stdout('C:\AMQSCLM\rpts\+QMNAME+.TSTCLUS.stdout.log') +
 stderr('C:\AMQSCLM\rpts\+QMNAME+.TSTCLUS.stderr.log')
```

| Definition      | Description                                                                                                                                                                                                                                                                                  |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>service</b>  | Specifies the service name. You can choose the service name.                                                                                                                                                                                                                                 |
| <b>descr</b>    | Specifies a textual description of the service.                                                                                                                                                                                                                                              |
| <b>control</b>  | Indicates that the service starts and stops at the same time as the queue manager.                                                                                                                                                                                                           |
| <b>servtype</b> | Indicates a server service object, meaning only one instance, can be executed at a time for this queue manager.                                                                                                                                                                              |
| <b>startcmd</b> | Specifies the location and name of the program.                                                                                                                                                                                                                                              |
| <b>startarg</b> | Specifies the arguments of the sample. Note the use of +QMNAME+. The name of the queue manager is automatically substituted.                                                                                                                                                                 |
| <b>stdout</b>   | The fully qualified file name to which standard output is redirected. The sample writes to this file only messages confirming that the sample has terminated. The sample does this because the standard error file has already closed in an earlier stage of the sample termination process. |
| <b>stderr</b>   | The fully qualified file name to which standard error output is redirected. The sample writes to the standard error file any error messages prior to termination of the sample.                                                                                                              |

### About this task

This task enables you to start and stop the cluster queue monitoring sample in different ways. It also enables you to run the sample in a mode that generates report files containing statistical information about the queues being monitored.

The sample program can be run by using the following command.

```
AMQSCLM -m QMgrName -c ClusterName (-q QNameMask| -f QListFile) -r MonitorQName
[-l ReportDir] [-t] [-u ActiveVal] [-i Interval] [-d] [-s] [-v]
```

| Argument  | Variable           | Further Information                           |
|-----------|--------------------|-----------------------------------------------|
| <b>-m</b> | <b>QMgrName</b>    | The queue manager to monitor.                 |
| <b>-c</b> | <b>ClusterName</b> | The cluster containing the queues to monitor. |



| Argument  | Variable            | Further Information                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-q</b> | <b>QNameMask</b>    | The queue, or queues, to monitor. A trailing * monitors all queues with names that match zero or more trailing characters.                                                                                                                                                                                                                                                                                                                                                          |
| <b>-f</b> | <b>QListFile</b>    | The full path and file name of a file containing a list of queue names or queue name masks to monitor. The file must contain one queue name/mask per line. You can specify <b>-q</b> or <b>-f</b> , but not both.                                                                                                                                                                                                                                                                   |
| <b>-r</b> | <b>MonitorQName</b> | The local queue being used exclusively by the sample.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>-l</b> | <b>ReportDir</b>    | The directory path in which to store logged information messages in a set of wrapping(For each queue manager and queue combination a fixed-size log file is generated that, when full, is overwritten. The logger always writes into the same file, and also keeps the two previous versions of the file.) report files.                                                                                                                                                            |
| <b>-t</b> |                     | (Optional) Enables the transfer of queued messages from inactive local queues to active queues. If not enabled, only new messages entering the cluster are dynamically routed to active instances of a queue.                                                                                                                                                                                                                                                                       |
| <b>-u</b> | <b>ActiveVal</b>    | (Optional) Automatically switches the <b>CLWLUSEQ</b> property of a monitored queue instance to ANY when it is inactive, and to the value of <b>ActiveVal</b> when active. <b>ActiveVal</b> can be LOCAL or QMGR. If this argument is not set in a system where putting applications connect to the same queue manager, or where message transfer is enabled, then monitored queues must have a <b>CLWLUSEQ</b> value of ANY, or QMGR with the queue manager having a value of ANY. |
| <b>-i</b> | <b>Interval</b>     | (Optional) The time interval in seconds, at which the monitor checks the queues. Default is 300 seconds (5 minutes).                                                                                                                                                                                                                                                                                                                                                                |
| <b>-d</b> |                     | (Optional) Enables additional diagnostic output. Debug output might be useful when initially configuring the system, or when working with the sample code.                                                                                                                                                                                                                                                                                                                          |
| <b>-s</b> |                     | (Optional) Enables minimal statistical output per interval.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>-v</b> |                     | (Optional) Log report information to standard out, in addition to the report files.                                                                                                                                                                                                                                                                                                                                                                                                 |

Argument list examples:

```
-m QMGR1 -c CLUS1 -f c:\QList.txt -r CLMQ -l c:\amqsc\m\rpts -s
-m QMGR2 -c CLUS1 -q ABC* -r CLMQ -l c:\amqsc\m\rpts -i 600
-m QMGR1 -c CLUSDEV -q QUEUE.* -r CLMQ -l c:\amqsc\m\rpts -t -u QMGR -d
```

Example queue list file:

```
Q1
QUEUE.*
ABC
ABD
```

## Procedure

1. Start the cluster queue monitoring sample. You can start the sample in one of the following ways:
  - Use a command prompt with the appropriate user authorizations.
  - Use the MQSC **START SERVICE** command, if the sample is configured as an IBM MQ service.

The argument list is the same in both cases. The sample does not start monitoring the queues for 10 seconds after the program is initialized. This delay allows consuming applications to connect to the monitored queues first, preventing unnecessary changes to the active state of the queue.

2. Stop the cluster queue monitoring sample. The sample automatically stops when the queue manager is stopped, stopping, quiescing, or if the connection to the queue manager is broken. There are ways to stop the sample without ending the queue manager:
  - Configure the local queue used exclusively by the sample to disable the Get function.
  - Send a message with a **CorrelId** of "STOP CLUSTER MONITOR\0\0\0\0", to the local queue used exclusively by the sample.
  - Terminate the sample process. This might result in the loss of non-persistent messages being transferred to active queues. It might also result in the local queue used by the sample being held open for a number of seconds after the termination. This situation prevents a new instance of the cluster queue monitoring sample from starting immediately.

If the sample has been started as an IBM MQ service, **STOP SERVICE** has no effect. It is possible to use one of the termination methods described as a configured **STOP SERVICE** mechanism in the queue manager.

## What to do next

Check the status of the sample.

If reporting is enabled, you can review the report files for status. Use the following command to review the most current report file:

```
QMgrName.ClusterName.RPT01.LOG
```

To review older report files, use the following commands:

```
QMgrName.ClusterName.RPT02.LOG
QMgrName.ClusterName.RPT03.LOG
```

Report files grow to a maximum size of approximately 1 MB. When the RPT01 file fills up, a new RPT01 file is created. The old RPT01 file is renamed to RPT02. RPT02 is renamed to RPT03. The old RPT03 is discarded.

The sample creates information messages in the following situations:

- at startup
- at termination

- when it marks a queue **ACTIVE** or **INACTIVE**
- when it requeues messages from an inactive queue to an active instance or instances

The sample creates an error message *CLMnnnnE* to report a problem that requires attention.

Every 30 minutes, the sample reports average processing time per polling interval, and elapsed processing time. This information is held in message CLM0045I.

When statistical messages are enabled **-s**, the sample reports the following statistical information about each queue check:

- Time taken to process the queues (in milliseconds)
- Number of queues checked
- Number of active/inactive changes made
- Number of messages transferred

This information is reported in message CLM0048I.

Report files might grow rapidly in debug mode, and quickly wrap. In this situation, the 1 MB size limit for individual files might be exceeded.

*AMQSCLM: Troubleshooting:*

The following sections contain information about scenarios that might be encountered while using the sample. Information about potential explanations for a scenario, and options on how to resolve it, are provided.

**Scenario: AMQSCLM is not starting**

**Potential explanation:** Incorrect syntax.

**Action:** Check standard error output for correct syntax

**Potential explanation:** Queue manager is not available.

**Action:** Check the report file for message ID CLM0010E.

**Potential explanation:** Cannot open or create report file or files.

**Action:** Check standard error output for error messages during initialization.

**Scenario: AMQSCLM is not changing a queue to ACTIVE or INACTIVE**

**Potential explanation:** The queue is not in the list of queues to be monitored

**Action:** Check the **-q** and **-f** parameter values.

**Potential explanation:** The queue is not a local queue in the correct cluster.

**Action:** Check that the queue is local and in the correct cluster.

**Potential explanation:** AMQSCLM is not running for this queue manager and cluster.

**Action:** Start AMQSCLM for the relevant queue manager and cluster.

**Potential explanation:** The queue is left INACTIVE, **CLWLPRTY** =0, because it has no consumers. Alternatively, it is left ACTIVE **CLWLPRTY** >=1, because it has at least 1 consumer.

**Action:** Check whether consuming applications are attached to the queue.

**Potential explanation:** The command server of the queue manager is not running.

**Action:** Check the report files for errors.

#### **Scenario: Messages are not being routed around INACTIVE queues**

**Potential explanation:** Messages are put directly to the queue manager that owns the inactive queue, and the **CLWLUSEQ** value of the queue is not ANY, and the **-u** argument is not being used for AMQSCLM.

**Action:** Check the **CLWLUSEQ** value of the relevant queue manager, or ensure that the **-u** argument is used for AMQSCLM.

**Potential explanation:** There are no active queues on any queue managers. Messages are evenly workload balanced across all inactive queues until a queue becomes active.

**Action:** Check the status of the queues on all queue managers.

**Potential explanation:** Messages are put to a different queue manager in the cluster to the one that owns the inactive queue, and the updated **CLWLPRTY** value of 0 is not propagated to the queue manager of the putting application.

**Action:** Check that the cluster channels between the monitored queue manager, and the full repository queue manager, are running. Check that the channels between the putting queue manager, and the full repository queue manager, are running. Check the error logs of the monitored, putting, and full repository queue managers.

**Potential explanation:** The remote queue instances are active (**CLWLPRTY**=1), but messages cannot be routed to those queue instances because the cluster sender channel from the local queue manager is not running.

**Action:** Check the status of the cluster sender channels from the local queue manager to the remote queue manager, or managers, with an active instance of the queue.

#### **Scenario: AMQSCLM is not transferring messages from an inactive queue**

**Potential explanation:** Message transfer is not enabled ( **-t** ).

**Action:** Ensure that message transfer is enabled ( **-t** ).

**Potential explanation:** The queue is not in the list of queues to be monitored.

**Action:** Check the **-q** and **-f** parameter values.

**Potential explanation:** AMQSCLM is not running for this, or other queue managers in the cluster, that own instances of the same queue.

**Action:** Start AMQSCLM.

**Potential explanation:** The queue has **CLWLUSEQ** = LOCAL or **CLWLUSEQ** = QMGR, and the **-u** argument is not set.

**Action:** Set the **-u** parameter, or change the queue, or queue manager configuration, to ANY.

**Potential explanation:** There are no active instances of the queue in the cluster.

**Action:** Check for instances of the queue with a **CLWLPRTY** value of 1, or greater.

**Potential explanation:** Remote queue instances have consumers ( **IPPROCS** >=1) but are inactive on those queue managers ( **CLWLPRTY** =0) because AMQSCLM is not monitoring those remote instances.

**Action:** Ensure that AMQSCLM is running on those queue managers, and / or the queue is in the list of queues to be monitored by checking the **-q** and **-f** parameter values.

**Potential explanation:** The remote queue instances are active ( **CLWLPRTY** =1), but are seen as inactive on the local queue manager ( **CLWLPRTY** =0). This situation is due to the updated **CLWLPRTY** value not being propagated to this queue manager.

**Action:** Ensure that the remote queue managers are connected to at least one of the full repository queue managers in the cluster. Ensure that the full repository queue managers are functioning correctly. Check that the channels between the full repository queue managers, and the monitored queue managers, are running.

**Potential explanation:** The messages are not committed, therefore they are not retrievable.

**Action:** Check that the sending application is functioning correctly.

**Potential explanation:** AMQSCLM does not have access to the local queue where messages are queued.

**Action:** Check whether AMQSCLM is running as a user with sufficient authorization to access the queue.

**Potential explanation:** The command server of the queue manager is not running.


**Action:** Start the command server of the queue manager.

**Potential explanation:** AMQSCLM encountered an error.

**Action:** Check the report files for errors.

**Potential explanation:** The remote queue instances are active (CLWLPRTY=1), but messages cannot be transferred to those queue instances because the cluster sender channel from the local queue manager is not running. This is often accompanied by a CLM0030W warning in the amqsclm report log.

**Action:** Check the status of the cluster sender channels from the local queue manager to the remote queue manager, or managers, with an active instance of the queue.

**Sample program for Connection Endpoint Lookup (CEPL):** 

IBM MQ Connection Endpoint Lookup sample provides a simple yet powerful exit module that offers IBM MQ users a way to retrieve connection definitions from an LDAP repository such as Tivoli Directory Server.

Tivoli Directory Server v6.3 Client must be installed in order to use CEPL.

A working knowledge of IBM MQ administration on the supported platforms is required to use this sample.


*Introduction:*    

Configure a global repository, for example, an LDAP (Lightweight Directory Access Protocol) directory, to store client connection definitions to aid maintenance and administration.

Using an IBM MQ Client application to establish a connection to a Queue Manager via a Client Connection Definition Table (CCDT).

The CCDT is created through the standard IBM MQ MQSC Administration interface. The user must be connected to a Queue Manager in order to create client connection definitions, even though the data contained within the definition is not restricted to the Queue Manager. The CCDT file generated must be

manually distributed among client machines and applications.

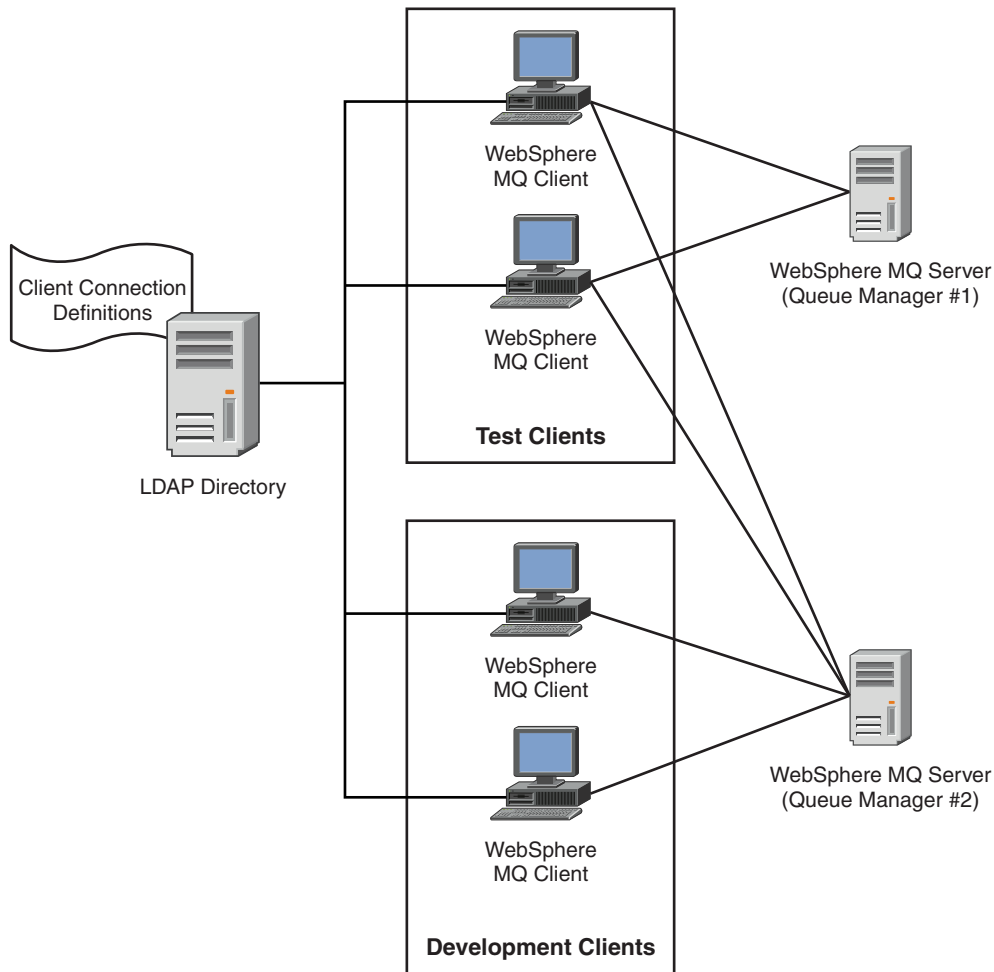


The CCDT file must be distributed to each IBM MQ client. Where thousands of clients can exist either locally or globally, it would soon become difficult to maintain and administer. A more flexible approach is needed to help ensure that each client has the correct client definitions available to them.

One such approach is to store the client connection definitions in a global repository such as an LDAP (Lightweight Directory Access Protocol) directory. An LDAP directory can also provide additional security, indexing, and search facilities, thereby allowing each client access to only those connection definitions pertaining to them.

The LDAP directory can be configured so that only specific definitions are available to certain user groups. For example, the Test Clients can access both Queue Manager #1 and #2, whereas the

Development Clients can access Queue Manager #2 only.




The exit module can look up an LDAP repository, for example, IBM Tivoli Directory Server, to retrieve channel definitions. Using those connection definitions, an IBM MQ client application can establish connection to a queue manager.

The exit module is a pre-connect exit module which enables channel definition to be obtained during the MQCONN/MQCONN call from an LDAP repository.

The exit module and schema might be implemented by:

- Customers who have already built a skill base using the existing CCDT file based technology and want to ease the administration and distribution costs.
- Existing customers who already employ their own propriety technology for distributing client connection definitions.
- New or existing customers who currently do not employ any type of client connection solution and want to use the features offered by IBM MQ.
- New or existing customers who want to directly use or tune their messaging model inline with any current LDAP business architecture.








Supported environments: 

Verify that you have a supported operating system and the relevant software before running the Connection Endpoint Lookup sample.




The sample program for IBM MQ Connection Endpoint Lookup requires the following software:


- IBM WebSphere MQ Version 7.0, or later
- Tivoli Directory Server V6.3 Client, or later

Operating Systems supported:

1.  Windows (7/8/2008/2012)
2.  Solaris (SPARC and x86-64)
3.  AIX
4.  Linux
  - RHEL v4 and v5 on System p
  - SUSE v9 and v10 on System p
  - RHEL v4 and v5 System x 32 bit and 64 bit
  - SUSE v9 and v10 System x 32 bit and 64 bit
5.  HP IA64.

**Note:** The sample is not available for the following platforms:

-  z/OS
-  IBM i
-  HP PA-RISC

Installation and configuration: 

Installing and configuring the exit module and connection endpoint schema.

### Installing exit module

During installation of IBM MQ, the exit module is installed under `tools/samples/c/preconnexit/bin`. For 32 bit platforms, the exit module must be copied to `exit/installation_name/` before it can be used. For 64 bit platforms, the exit module must be copied to `exit64/installation_name/` before it can be used.

### Installing Connection Endpoint schema

The exit uses the Connection Endpoint schema, `ibm-amq.schema`. The schema file must be imported into any LDAP server before the exit can be used. After importing the schema, values for the attributes must be added.

Here is an example for importing Connection Endpoint schema. The example assumes IBM Tivoli Directory Server (ITDS) is being used.

- Ensure that IBM Tivoli Directory Server is running, then copy or FTP the `ibm-amq.schema` file to the ITDS server.
- On the ITDS server, enter the following command to install the schema into ITDS store, where `LDAP ID` and `LDAP password` are the root DN and password for the LDAP server:

```
ldapadd -D "LDAP ID" -w "LDAP password" -f ibm-amq.schema
```

- In a command window, enter the following command or use a third party tool to browse the schema for verification:

```
ldapsearch objectclass=ibm-amqClientConnection
```

Refer to your LDAP Server documentation for further details on importing the schema file.

## Configuration

A new section called PreConnect must be added to the client configuration file, for example `mqclient.ini`. The PreConnect section contains the following keywords:


**Module** : The name of the module containing the API exit code. If this field contains the full path of the module, it is used as is. Otherwise the `exit` or `exit64` folder in the IBM MQ installation is searched.

**Function** : The name of the functional entry point into the library that contains the PreConnect exit code. The function definition adheres to the `MQ_PRECONNECT_EXIT` prototype.

**Data** : URI of the LDAP repository containing channel definitions.

The following snippet is an example of the changes required in the `mqclient.ini` file.

```
PreConnect:
Module=amq1celp
Function=PreConnectExit
Data=ldap://myLDAPServer.com:389/cn=wmq,ou=ibm,ou=com
Sequence=1
```

Overview of the exit and schema: 

Syntax, and parameters used to establish a connection to a queue manager.

IBM WebSphere MQ Version 7.0.1, Fix Pack 4 defines the following syntax for an entry point in an exit module.

```
void MQENTRY MQ_PRECONNECT_EXIT (PMQNX pExitParms
 , PMQCHAR pQMgrName
 , PPMQCNO ppConnectOpts
 , PMQLONG pCompCode
 , PMQLONG pReason)
```

During the `MQCONN/X` call execution, IBM MQ C Client loads the exit module containing an implementation of the function syntax. It then invokes an exit function to retrieve channel definitions. The retrieved channel definitions are then used to establish connection to a queue manager.

## Parameters

### pExitParms

Type: `PMQNX` input/output

The PreConnection exit parameter structure. The structure is allocated and maintained by the caller of the exit.

```
struct tagMQNX
{
 MQCHAR4 StrucId; /* Structure identifier */
 MQLONG Version; /* Structure version number */
 MQLONG ExitId; /* Type of exit */
 MQLONG ExitReason; /* Reason for invoking exit */
 MQLONG ExitResponse; /* Response from exit */
 MQLONG ExitResponse2; /* Secondary response from exit */
 MQLONG Feedback; /* Feedback code (reserved) */
 MQLONG ExitDataLength; /* Exit data length */
}
```

```

 PMQCHAR pExitDataPtr; /* Exit data */
 MQPTR pExitUserAreaPtr; /* Exit user area */
 PMQCD * ppMQCDArrayPtr; /* Array of pointers to MQCDs */
 MQLONG MQCDArrayCount; /* Number of entries found */
 MQLONG MaxMQCDVersion; /* Maximum MQCD version */
};

```

### **pQMgrName**

Type: PMQCHAR input/output

Name of the queue manager. On input, this parameter is the filter string supplied to the MQCONN API call through the **QMgrName** parameter. This field might be blank, explicit, or contain certain wildcard characters. The field is changed by the exit. The parameter is NULL when the exit is called with MQXR\_TERM.

### **ppConnectOpts**

Type: ppConnectOpts input/output

Options that control the action of MQCONN. This is a pointer to an MQCNO connection options structure that controls the action of the MQCONN API call. The parameter is NULL when the exit is called with MQXR\_TERM. The MQI client always provides an MQCNO structure to the exit, even if it was not originally provided by the application. If an application provides an MQCNO structure, the client makes a duplicate to pass it to the exit where it is modified. The client retains the ownership of the MQCNO. An MQCD referenced through the MQCNO takes precedence over any connection definition provided through the array. The client uses the MQCNO structure to connect to the queue manager and the others are ignored.

### **pCompCode**

Type: PMQLONG input/output

Completion code. Pointer to an MQLONG that receives the exits completion code. It must be one of the following values:

- MQCC\_OK - Successful completion
- MQCC\_WARNING - Warning (partial completion)
- MQCC\_FAILED - Call failed

### **pReason**

Type: PMQLONG input/output

Reason qualifying pCompCode. Pointer to an MQLONG that receives the exit reason code. If the completion code is MQCC\_OK, the only valid value is: MQRC\_NONE - (0, x'000') No reason to report.

If the completion code is MQCC\_FAILED or MQCC\_WARNING, the exit function can set the reason code field to any valid MQRC\_\* value.

The exit uses the following data structure for context information.

### MQNLDAPECTX

The MQNLDAPECTX structure has the following C prototype.

```
typedef struct tagMQNLDAPECTX MQNLDAPECTX;
typedef MQNLDAPECTX MQPOINTER PMQNLDAPECTX;

struct tagMQNLDAPECTX
{
 MQCHAR4 StrucId; /* Structure identifier */
 MQLONG Version; /* Structure version number */
 LDAP * objectDirectory /* LDAP Instance */
 MQLONG ldapVersion; /* Which LDAP version to use? */
 MQLONG port; /* Port number for LDAP server*/
 MQLONG sizeLimit; /* Size limit */
 MQBOOL ssl; /* SSL enabled? */
 MQCHAR * host; /* Hostname of LDAP server */
 MQCHAR * password; /* Password of LDAP server */
 MQCHAR * searchFilter; /* LDAP search filter */
 MQCHAR * baseDN; /* Base Distinguished Name */
 MQCHAR * charSet; /* Character set */
};
```

Sample code for building the connection endpoint lookup exit:

AIX

Linux

Solaris

Windows

You can use the sample code snippets for compiling the source on AIX, Linux, Solaris, and Windows.

### Compiling source

You can compile the source with any LDAP client libraries, for example, IBM Tivoli Directory Server V6.3 Client libraries. This documentation assumes that you are using Tivoli Directory Server V6.3 client libraries.

**Note:** The pre-connect exit library is supported with the following LDAP servers:

- IBM Tivoli Directory Server V6.3
- Novell eDirectory V8.2

The following code snippets describe how to compile the exits:

Windows

#### Compiling the exit on the Windows platform

You can use the following snippet for compiling the exit source:

```
CC=c1.exe
LL=link.exe
CCARGS=/c /I. /DWIN32 /W3 /DNDEBUG /EHsc /D_CRT_SECURE_NO_DEPRECATED /Z1

The libraries to include
LDLIBS=ws2_32.lib Advapi32.lib libibmldapstatic.lib libibmldapdbgstatic.lib \
kernel32.lib user32.lib gdi32.lib winspool.lib comdlg32.lib advapi32.lib \
shell32.lib ole32.lib oleaut32.lib uuid.lib odbc32.lib odbccp32.lib msvcrt.lib

OBSJ=amqlcel0.obj

all: amqlcelp.dll

amqlcelp.dll: $(OBSJ)
 $(LL) /OUT:amqlcelp.dll /INCREMENTAL /NOLOGO /DLL /SUBSYSTEM:WINDOWS /MACHINE: X86 \
 /DEF:amqlcelp.def $(OBSJ) $(LDLIBS) /NODEFAULTLIB:msvcrt.lib
```

```
The exit source
amqlce10.obj: amqlce10.c
$(CC) $(CCARGS) $*.c
```

**Note:** If you are using the IBM Tivoli Directory Server V6.3 Client libraries that are compiled with Microsoft Visual Studio 2003 compiler, you might get warnings while you are compiling the IBM Tivoli Directory Server V6.3 Client libraries with Microsoft Visual Studio 2012, or later, compiler.

### AIX Linux Solaris **Compiling the exit on AIX, Linux, or Solaris**

The following code snippet is for compiling the exit source on Linux. Some compiler options might differ on AIX or Solaris.

```
#Make file to build exit
CC=gcc
```

```
MQML=/opt/mqm/lib
MQMI=/opt/mqm/inc
TDSI=/opt/ibm/ldap/V6.3/include
XFLAG=-m32
```

```
TDSL=/opt/ibm/ldap/V6.3/lib
```

IBM Tivoli Directory Server ships both static and dynamic link libraries, but you can use only one type of library. This script assumes that you are using the static libraries.

```
#Use static libraries.
LDLIBS=-L$(TDSL) -libibmldapstatic
```

```
CFLAGS=-I. -I$(MQMI) -I$(TDSI)
```

```
all:amqlcepl
```

```
amqlcepl: amqlce10.c
$(CC) -o cepl amqlce10.c -shared -fPIC $(XFLAG) $(CFLAGS) $(LDLIBS)
```

Invocation of the PreConnect exit module: ULW

The PreConnect exit module can be invoked with three different reason codes: the MQXR\_INIT reason code for initializing and establishing a connection to an LDAP server, the MQXR\_PRECONNECT reason code for retrieving channel definitions from an LDAP server, or the MQXR\_TERM reason code when the exit is to be cleaned.

#### **MQXR\_INIT**

The exit is invoked with MQXR\_INIT reason code for initializing and establishing a connection to an LDAP server.

Before the MQXR\_INIT call, the pExitDataPtr field of the MQNXP structure is populated with the Data attribute from the PreConnect stanza within the mqclient.ini file (that is, the LDAP).

An LDAP URL consists of at least the protocol, host name, port number, and base DN for the search. The exit parses the LDAP URL contained within the pExitDataPtr field, allocates an MQNLDAPCTX LDAP Lookup Context structure and populates it accordingly. The address of this structure is stored in the pExitUserAreaPtr field. Failure to correctly parse the LDAP URL results in the error MQCC\_FAILED.

At this point, the exit connects and binds to the LDAP server using the MQNLDAPCTX parameters. The resulting LDAP API handles are also stored within this structure.

#### **MQXR\_PRECONNECT**

The exit module is invoked with the MQXR\_PRECONNECT reason code for retrieving channel definitions from an LDAP server.

The exit searches the LDAP server for channel definitions matching the given filter. If the **QMgrNameparameter** contains a specific queue manager name, the search returns all channel definitions for which the **ibm-amqQueueManagerName** LDAP attribute value matches with the given queue manager name.

If the **QMgrName** parameter is '\*' or '(blank)', then the search returns all channel definitions for which the **ibm-amqIsClientDefault Connection** endpoint attribute is set to TRUE.


After a successful search, the exit prepares one or an array of MQCD definitions and returns back to the caller.

## MQXR\_TERM

The exit is invoked with this reason code when the exit is to be cleaned. During this cleaning, the exit disconnects from the LDAP server, and releases all the memory allocated and maintained by the exit, including the MQNLDAPCTX structure, the pointer array, and every MQCD it references. Any other fields are set to the default values. The **pQMgrName** and **ppConnectOpts** exit parameters are unused during an exit with the MQXR\_TERM reason code and may be NULL.

### Related information:

PreConnect stanza of the client configuration file

LDAP schemas: 

Client connection data is stored in a global repository called the LDAP (Lightweight Directory Access Protocol) directory. An IBM MQ client uses an LDAP directory to obtain the connection definitions. The structure of the IBM MQ client connection definitions within the LDAP directory is known as the LDAP schema. An LDAP schema is the collection of attribute type definitions, object class definitions, and other information which a server uses to determine whether a filter or attribute value assertion matches against the attributes of an entry, and whether to permit, add, and modify operations.

### Storing data in the LDAP directory

The client connection definitions are located under a specific branch within the directory tree known as the connection point. Like all other nodes within an LDAP directory, the connection point has a Distinguished Name (DN) associated with it. You can use this node as the starting point for any queries you make on the directory. Use filtering when querying the LDAP directory to return a subset of client connection definitions. You can restrict access to sub-trees based on permissions granted in other parts of the directory tree - for example, to users, departments, or groups.

### Defining your own attributes and classes

Store the client channel definition by modifying the LDAP schema. All LDAP data definitions require objects and attributes. The objects and attributes are identified by an object identifier (OID) number which uniquely identifies the object or attribute. All classes within an LDAP schema inherit either directly or indirectly from the top object. The client channel definition object contains the attributes of the top object. All LDAP data definitions require objects and attributes:

- Object definitions are collections of LDAP attributes.
- Attributes are LDAP data types.

The description of each attribute and how they map to the normal IBM MQ properties are described in LDAP attributes.

LDAP attributes defined are specific to IBM MQ and maps directly to the client connection properties.

### IBM MQ Client Channel Directory String Attributes

The character string attributes with their mapping to IBM MQ properties are listed in the following table. The attributes can hold values of directoryString (UTF-8 encoded Unicode, that is, a variable byte encoding system that includes IA5/ASCII as a subset) syntax. The syntax is specified by its object identification number (OID).

Table 79. IBM MQ client channel directory string attributes

| LDAP Attribute                | Description                                                                                                                        | IBM MQ Property |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| CN                            | The common name consisting of the channel name and the defining queue manager name.                                                |                 |
| ibm-amqChannelName            | The name of the channel definition.                                                                                                | CHANNEL         |
| ibm-amqConnectionName         | The communication connection identifier.                                                                                           | CONNNAME        |
| ibm-amqDescription            | The channel description.                                                                                                           | DESCR           |
| ibm-amqLocalAddress           | The local communication address of the channel.                                                                                    | LOCLADDR        |
| ibm-amqModeName               | The LU 6.2 mode name.                                                                                                              | MODENAME        |
| ibm-amqPassword               | The password that can be used.                                                                                                     | PASSWORD        |
| ibm-amqQueueManagerName       | The name of the queue manager or queue manager group to which an IBM MQ client application can request connection.                 | QMNAME          |
| ibm-amqSecurityExitUserData   | The user data that is passed to the security exit.                                                                                 | SCYDATA         |
| ibm-amqSecurityExitName       | The name of the exit program to be run by the channel security exit.                                                               | SCYEXIT         |
| ibm-amqSslCipherSpec          | A single CipherSpec for a TLS connection.                                                                                          | SSLCIPH         |
| ibm-amqSslPeerName            | Checks the Distinguished Name (DN) of the certificate from the peer queue manager or client at the other end of an IBM MQ channel. | SSLPEER         |
| ibm-amqTransactionProgramName | The transaction program name.                                                                                                      | TPNAME          |
| ibm-amqUserID                 | The user ID to be used by the MCA when attempting to initiate a secure SNA session with a remote MCA.                              | USERID          |

### IBM MQ client connection integer attributes

The attributes with predefined values (for example, an enumerated type) are stored as standard integers. These values are stored in the LDAP directory as integer values, and not by using the associated constant name.

Table 80. IBM MQ client channel directory integer attributes

| LDAP attribute             | Description                                                                                                                                 | IBM MQ Property |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| ibm-amqConnectionAffinity  | Determines whether client applications, which connect multiple times through the same queue manager name, use the same client channel.      | AFFINITY        |
| ibm-amqClientChannelWeight | A weighting to influence which client connection channel definition is used.                                                                | CLNTWGHT        |
| ibm-amqHeartBeatInterval   | The approximate time between heartbeat flows that are to be passed from a sending MCA when there are no messages on the transmission queue. | HBINT           |
| ibm-amqKeepAliveInterval   | A time-out value for a channel.                                                                                                             | KAINT           |

Table 80. IBM MQ client channel directory integer attributes (continued)

| LDAP attribute              | Description                                                                  | IBM MQ Property |
|-----------------------------|------------------------------------------------------------------------------|-----------------|
| ibm-amqMaximumMessageLength | The maximum length of a message that can be transmitted on the channel.      | MAXMSGL         |
| ibm-amqSharingConversations | The maximum number of conversations that share each TCP/IP channel instance. | SHARECNV        |
| ibm-amqTransportType        | The transport type to be used.                                               | TRPTYPE         |

#### IBM MQ client channel boolean attribute

This Boolean attribute is not mapped to any IBM MQ property. The syntax of this attribute indicates a boolean value.

Table 81. IBM MQ client channel boolean attribute

| LDAP attribute         | Description                                                                                                                                              |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| ibm-amqIsClientDefault | This boolean attribute is defined to resolve the problem of searching entries whose <code>ibm-amqQueueManagerName</code> attribute has not been defined. |


#### IBM MQ client channel list attributes

IBM MQ properties are stored as single-value, comma-separated list attribute within the LDAP directory. The attributes are defined in the same manner as the other directory string attributes. The list attributes along with their mapping to the IBM MQ properties are described in the following table.

Table 82. IBM MQ client channel list attributes

| LDAP attribute             | Description                                                                   | IBM MQ Property |
|----------------------------|-------------------------------------------------------------------------------|-----------------|
| ibm-amqHeaderCompression   | A list of header data compression techniques supported by the channel.        | COMPHDR         |
| ibm-amqMessageCompression  | A list of message data compression techniques supported by the channel.       | COMPMSG         |
| ibm-amqSendExitUserData    | The user data that is passed to the send exit.                                | SENDDATA        |
| ibm-amqSendExitUserName    | The name of the exit program to be run by the channel send exit.              | SENDEXIT        |
| ibm-amqReceiveExitUserData | The user data that is passed to the receive exit.                             | RCVDATA         |
| ibm-amqReceiveExitName     | The name of the user exit program to be run by the channel receive user exit. | RCVEXIT         |



*Common Name:* 

The common name (CN) consists of the channel name and the defining queue manager name.

It is a preexisting attribute.

The format of the CN is:

CN=CHANNEL\_NAME(DEFINING\_Q\_MGR\_NAME)

For example:

CN=TC1(QM\_T1)

You can specify only one value for this attribute.


This attribute is a string attribute and the values are not case-sensitive. The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter, using a substring (for example, CN=jim\* where CN is an attribute) and contains one or more wildcards.

*ibm-amqChannelName:* 

This attribute specifies the name of the channel definition.

This attribute has a single string value with a maximum of 20 characters that are not case-sensitive. It is not a pre-existing attribute.


The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter, using a substring and contains one or more wildcards.

*ibm-amqDescription:* 

This LDAP attribute provides the channel description.

This attribute has a single string value with a maximum of 64 bytes, which are not case-sensitive. It is not a pre-existing attribute.


The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqConnectionName*: 

This LDAP attribute is the communications connection identifier. It specifies the particular communications links to be used by this channel.

This attribute has a single string value with a maximum of 264 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqLocalAddress*: 

This attribute specifies the local communications address for the channel.

This attribute has a single string value with a maximum of 48 characters, which are not case-sensitive. It is not a pre-existing attribute.


The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqModeName*: 

This attribute is for use with LU 6.2 connections. It gives extra definition for the session characteristics of the connection when a communication session allocation is performed.

This attribute has a single string value of exactly 8 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqPassword*: 

This LDAP attribute specifies a password that can be used by the MCA when attempting to initiate a secure LU 6.2 session with a remote MCA.

This attribute has a single integer value with a maximum of 12 digits. It is not a pre-existing attribute.

*ibm-amqQueueManagerName*: 

This attribute specifies the name of the queue manager or queue manager group to which an IBM MQ client application can request connection.

This attribute has a single string value with a maximum of 48 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

**Related reference:**

“ibm-amqIsClientDefault” on page 565

This Boolean attribute solves the problem of searching entries where the `ibm-amqQueueManagerName` attribute has not been defined.

*ibm-amqSecurityExitUserData:* 

This LDAP attribute specifies user data that is passed to the security exit.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqSecurityExitName:* 

This LDAP attribute specifies the name of the exit program to be run by the channel security exit.

Leave blank if no channel security exit is in effect.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. This attribute is not a pre-existing one.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqSslCipherSpec:* 

This LDAP attribute specifies a single CipherSpec for a TLS connection.

This attribute has a single string value with a maximum of 32 characters, which are not case-sensitive. It is not a pre-existing attribute.


The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqSslPeerName:* 

This LDAP attribute is used to check the Distinguished Name (DN) of the certificate from the peer queue manager or client at the other end of an IBM MQ channel.

This LDAP attribute has a single string value with a maximum of 1024 bytes, which are not case-sensitive. It is not a pre-existing one.


The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqTransactionProgramName:* 

This LDAP attribute specifies the transaction program name. It is for use with LU 6.2 connections.

This attribute has a single string value with a maximum of 64 characters, which are not case-sensitive. It is not a pre-existing one.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqUserID:* 

This LDAP attribute specifies the user ID to be used by the MCA when attempting to initiate a secure SNA session with a remote MCA.

This attribute has a single string value of exactly 12 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

*ibm-amqConnectionAffinity:* 

This LDAP attribute specifies whether client applications, which connect multiple times using the same queue manager name, use the same client channel.

This attribute has a single integer value. It is not a pre-existing attribute.

*ibm-amqClientChannelWeight:* 

This LDAP attribute specifies a weighting that influences which client connection channel definition is used.


The client channel weighting attribute is used to bias the selection of client channel definitions when more than one suitable definition is available.

This attribute has a single integer value. It is not a pre-existing attribute.

*ibm-amqHeartBeatInterval:* 

This LDAP attribute specifies the approximate time between heartbeat flows that are to be passed from a sending MCA when there are no messages on the transmission queue.

This attribute has a single integer value. It is not a pre-existing attribute. The default value is 1. The default is set in the current MQSERVER environment variable operation.

*ibm-amqKeepAliveInterval:* 

This LDAP attribute is used to specify a time-out value for a channel.


The value of this attribute is passed to the communications stack specifying the keepalive timing for the channel. You can use this to specify a different keepalive value for each channel.

This attribute has a single integer value. It is not a pre-existing attribute.

*ibm-amqMaximumMessageLength:* 

This LDAP attribute specifies the maximum length of a message that can be transmitted on the channel.

The default value of this attribute is 104857600 as per the current MQSERVER environment variable operation. This attribute has a single integer value and it is not a pre-existing attribute.

*ibm-amqSharingConversations:* 

This LDAP attribute specifies the maximum number of conversations that share each TCP/IP channel instance.

This attribute has a single integer value. This attribute is not a pre-existing attribute.

*ibm-amqTransportType:* 

This LDAP attribute specifies the transport type to be used.

This attribute has a single integer value. It is not a pre-existing attribute.

*ibm-amqIsClientDefault:* 

This Boolean attribute solves the problem of searching entries where the *ibm-amqQueueManagerName* attribute has not been defined.

Preconnect exit modules generally search the LDAP servers with the value of the *ibm-amqQueueManagerName* attribute as the search criteria. Such a query would return all entries where the *ibm-amqQueueManagerName* attribute value matches the name of the queue manager specified on the MQCONN/X call. However when using the client channel definition tables (CCDT), you can either set the queue manager name on a MQCONN/X call as blank or prefix the name with an asterisk (\*). If the name of the queue manager is blank, the client connects to the default queue manager. If the name is prefixed with an asterisk (\*) to the queue manager, then the client connects any queue manager.

Similarly, the *ibm-amqQueueManagerName* attribute in an entry can be left undefined. In this case, it is expected that the client using this endpoint information can connect to any queue manager. For example, an entry contains the following lines:

```
ibm-amqChannelName = "CHANNEL1"
ibm-amqConnectionName = myhost(1414)
```

In this example, the client attempts to connect to the specified queue manager running on myhost.

However in LDAP Servers, a search is not made on an attribute value that has not been defined. For example, if an entry contains the connection information except *ibm-amqQueueManagerName*, then the search results would not include this entry. To overcome this problem, you can set *ibm-amqIsClientDefault*. This is a Boolean attribute and is assumed to have a value of FALSE if not defined.


For entries where the `ibm-amqQueueManagerName` has not been defined and are expected to be part of the search, set `ibm-amqIsClientDefault` to `TRUE`. When a blank or asterisk (\*) is specified as the queue manager name in a call to `MQCONN/X`, the preconnect exit searches the LDAP sever for all entries where `ibm-amqIsClientDefault` attribute value is set to `TRUE`.

**Note:** Do not set or define `ibm-amqQueueManagerName` attribute if `ibm-amqIsClientDefault` is set to `TRUE`.

**Related reference:**

“`ibm-amqQueueManagerName`” on page 562

This attribute specifies the name of the queue manager or queue manager group to which an IBM MQ client application can request connection.

*ibm-amqHeaderCompression:* 

This LDAP attribute is a list of header data compression techniques supported by the channel.

The maximum size of this attribute is of 48 characters. It is not a pre-existing attribute.

You can specify only one value for this attribute.

This list attribute is specified as directory strings using a comma-separated format. For example, the value specified for **`ibm-amqHeaderCompression`** is `0` which is mapped to `NONE`. Any values that exceed the maximum allowed limit are ignored by the client. For example, `ibm-amqHeaderCompression` contains a maximum of 2 integers in the list.

*ibm-amqMessageCompression:* 

This LDAP attribute is a list of message data compression techniques supported by the channel.

The maximum size of this attribute is of 48 characters. It is not a pre-existing attribute.

This attribute does not support multiple values.

This list attribute is specified as directory strings using a comma-separated format. For example, the value specified for this attribute is `1,2,4`, which maps to the underlying compression sequence `RLE`, `ZLIBFAST`, and `ZLIBHIGH`.

Any values that exceed the maximum allowed limit are ignored by the client. For example, `ibm-amqMessageCompression` contains a maximum of 16 integers in the list.


*ibm-amqSendExitUserData:* 

This LDAP attribute specifies user data that is passed to the send exit.

This LDAP attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

**Note:** **ibm-amqSendExitName** and **ibm-amqSendExitUserData** need to be synchronized in pairs . The user data should be synchronized with the exit name. So if one is specified, the other also has to be symmetrically specified, even if it contains no data.


*ibm-amqSendExitName:* 

This LDAP attribute specifies the name of the exit program to be run by the channel send exit.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

**Note:** **ibm-amqSendExitName** and **ibm-amqSendExitUserData** must be synchronized in pairs. The user data must be synchronized with the exit name. So if one is specified, the other also must be symmetrically specified even if it contains no data.

*ibm-amqReceiveExitUserData:* 

This LDAP attribute specifies user data that is passed to the receive exit.

You can run a sequence of receive exits. The string of user data for a series of exits is separated by a comma, spaces, or both.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

**Note:** **ibm-amqReceiveExitName** and **ibm-amqReceiveExitUserData** must be synchronized in pairs. The user data must be synchronized with the exit name. So if one is specified, the other also must be symmetrically specified even if it contains no data.

*ibm-amqReceiveExitName*: 

This LDAP attribute specifies the name of the user exit program to be run by the channel receive user exit.

This attribute is a list of names of programs that are to be run in succession. Leave blank, if no channel receive user exit is in effect.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

**Note:** *ibm-amqReceiveExitName* and *ibm-amqReceiveExitUserData* must be synchronized in pairs. The user data must be synchronized with the exit name. So if one is specified, the other must also be symmetrically specified, even if it contains no data.

## Using the sample programs for z/OS



The sample procedural applications that are delivered with IBM MQ for z/OS demonstrate typical uses of the Message Queue Interface (MQI).

### About this task

IBM MQ for z/OS also provides sample data-conversion exits, described in “Writing data-conversion exits” on page 364.

All the sample applications are supplied in source form; several are also supplied in executable form. The source modules include pseudocode that describes the program logic.

**Note:** Although some of the sample applications have basic panel-driven interfaces, they do not aim to demonstrate how to design the look and feel of your applications. For more information about how to design panel-driven interfaces for non-programmable terminals, see the *SAA Common User Access: Basic Interface Design Guide* (SC26-4583) and its addendum (GG22-9508). These provide guidelines to help you to design applications that are consistent both within the application and across other applications.

### Procedure


Use the following links to find out more about the sample programs:

- “Features demonstrated in the sample applications for z/OS” on page 569
- “Preparing and running sample applications for the batch environment on z/OS” on page 576
- “Preparing sample applications for the TSO environment on z/OS” on page 579
- “Preparing the sample applications for the CICS environment on z/OS” on page 581
- “Preparing the sample application for the IMS environment on z/OS” on page 584
- “The Put samples on z/OS” on page 585
- “The Get samples on z/OS” on page 587
- “The Browse sample on z/OS” on page 590
- “The Print Message sample on z/OS” on page 592
- “The Queue Attributes sample on z/OS” on page 596
- “The Mail Manager sample on z/OS” on page 597



- “The Credit Check sample on z/OS” on page 604
- “The Message Handler sample on z/OS” on page 617
- “The Asynchronous Put sample on z/OS” on page 622
- “The Batch Asynchronous Consumption sample on z/OS” on page 623
- “The CICS Asynchronous Consumption and Publish/Subscribe sample on z/OS” on page 624
- “The Publish/Subscribe sample on z/OS” on page 627
- “The Set and Inquire message property sample on z/OS” on page 629


**Related tasks:**

 “Using the sample programs on Multiplatforms” on page 457

These sample procedural programs are delivered with the product. The samples are written in C and COBOL, and demonstrate typical uses of the Message Queue Interface (MQI).

**Features demonstrated in the sample applications for z/OS:** 

This section summarizes the MQI features demonstrated in each of the sample applications, shows the programming languages in which each sample is written, and the environment in which each sample runs.


*Put samples on z/OS:* 

The Put samples demonstrate how to put messages on a queue using the MQPUT call.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQPUT
- MQCLOSE
- MQDISC

The program is delivered in COBOL and C, and runs in the batch and CICS environment. See Table 85 on page 577 for the batch application and Table 92 on page 582 for the CICS application.


*Get samples on z/OS:* 

The Get samples demonstrate how to get messages from a queue using the MQGET call.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQGET
- MQCLOSE
- MQDISC

The program is delivered in COBOL and C, and runs in the batch and CICS environment. See Table 85 on page 577 for the batch application and Table 92 on page 582 for the CICS application.


*Browse sample on z/OS:* 

The Browse sample demonstrates how to use the Browse option to find a message, print it, then step through the messages on a queue.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQGET for browsing messages
- MQCLOSE
- MQDISC

The program is delivered in the COBOL, assembler, PL/I, and C languages. The application runs in the batch environment. See Table 86 on page 577 for the batch application.

*Print Message sample on z/OS:* 


The Print Message sample demonstrates how to remove a message from a queue and print the data in the message, together with all the fields of its message descriptor. It can, optionally, display all of the message properties associated with each message.

By removing comment characters from two lines in the source module, you can change the program so that it browses, rather than removes, the messages on a queue. This program can usefully be used for diagnosing problems with an application that is putting messages on a queue.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQGET for removing messages from a queue (with an option to browse)
- MQCLOSE
- MQDISC
- MQCRTMH
- MQDLTMH
- MQINQMP

The program is delivered in the C language. The application runs in the batch environment. See Table 87 on page 578 for the batch application.

*Queue Attributes sample on z/OS:* 

The Queue Attributes sample demonstrates how to inquire about and set the values of IBM MQ for z/OS object attributes.

The application uses these MQI calls:

- MQOPEN
- MQINQ
- MQSET
- MQCLOSE

The program is delivered in the COBOL, assembler, and C languages. The application runs in the CICS environment. See Table 93 on page 582 for the CICS application.

*Mail Manager sample on z/OS:* 

Considerations to note when using Mail Manager sample.

The Mail Manager sample demonstrates these techniques:

- Using alias queues
- Using a model queue to create a temporary dynamic queue
- Using reply-to queues
- Using synchpoints in the CICS and batch environments
- Sending commands to the system-command input queue
- Testing return codes
- Sending messages to remote queue managers, both by using a local definition of a remote queue and by putting messages directly on a named queue at a remote queue manager

The application uses these MQI calls:


- MQCONN
- MQOPEN
- MQPUT1
- MQGET
- MQINQ
- MQCMIT
- MQCLOSE
- MQDISC

Three versions of the application are provided:

- A CICS application written in COBOL
- A TSO application written in COBOL
- A TSO application written in C

The TSO applications use the IBM MQ for z/OS batch adapter and include some ISPF panels.

See Table 90 on page 579 for the TSO application, and Table 94 on page 582 for the CICS application.

*Credit Check sample on z/OS:* 

This information contains points to consider when using Credit Check sample.

The Credit Check sample is a suite of programs that demonstrates these techniques:

- Developing an application that runs in more than one environment
- Using a model queue to create a temporary dynamic queue
- Using a correlation identifier
- Setting and passing context information
- Using message priority and persistence
- Starting programs by using triggering
- Using reply-to queues
- Using alias queues
- Using a dead-letter queue
- Using a namelist
- Testing return codes

The application uses these MQI calls:

- MQOPEN
- MQPUT
- MQPUT1
- MQGET for browsing and getting messages, using the wait and signal options, and for getting a specific message
- MQINQ
- MQSET
- MQCLOSE

The sample can run as a stand-alone CICS application. However, to demonstrate how to design a message queuing application that uses the facilities provided by both the CICS and IMS environments, one module is also supplied as an IMS batch message processing program.

The CICS programs are delivered in C and COBOL. The single IMS program is delivered in C.

See Table 95 on page 583 for the CICS application, and Table 97 on page 584 for the IMS application.


*The Message Handler sample on z/OS:* 

The Message Handler sample allows you to browse, forward, and delete messages on a queue.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQINQ
- MQPUT1
- MQCMIT
- MQBACK
- MQGET
- MQCLOSE
- MQDISC

The program is delivered in C and COBOL programming languages. The application runs under TSO. See Table 91 on page 580 for the TSO application.

*Distributed queuing exit samples on z/OS:* 

A table of source programs of Distributed queuing exit samples.

The names of the source programs of the distributed queuing exit samples are listed in the following table:

*Table 83. Source for the distributed queuing exit samples*

| Member name | For language | Description    | Supplied in library |
|-------------|--------------|----------------|---------------------|
| CSQ4BAX0    | Assembler    | Source program | SCSQASMS            |
| CSQ4BCX1    | C            | Source program | SCSQ37S             |
| CSQ4BCX2    | C            | Source program | SCSQ37S             |
| CSQ4BCX4    | C            | Source program | SCSQ37S             |

**Note:** The source programs are link-edited with CSQXSTUB.

*Data-conversion exit samples on z/OS:* 

A skeleton is provided for a data-conversion exit routine, and a sample is shipped with IBM MQ illustrating the MQXCNVC call.

The names of the source programs of the data-conversion exit samples are listed in the following table:

*Table 84. Source for the data conversion exit samples (assembler language only)*

| Member name | Description    | Supplied in library |
|-------------|----------------|---------------------|
| CSQ4BAX8    | Source program | SCSQASMS            |
| CSQ4BAX9    | Source program | SCSQASMS            |
| CSQ4CAX9    | Source program | SCSQASMS            |

**Note:** The source programs are link-edited with CSQASTUB.

See “Writing data-conversion exits” on page 364 for more information.

*Publish/Subscribe samples on z/OS:* 

The Publish/Subscribe sample programs demonstrate the use of the publish and subscribe features in IBM MQ.

There are four C and two COBOL programming language sample programs demonstrating how to program to the IBM MQ Publish/Subscribe interface.

The applications use these MQI calls:

- MQCONN
- MQOPEN
- MQPUT
- MQSUB
- MQGET

- MQCLOSE
- MQDISC
- MQCRTMH
- MQDLTMH
- MQINQMP

The Public/Subscribe sample programs are delivered in the C and COBOL programming languages. The sample applications run in the batch environment. See Publish/Subscribe samples for the batch applications.

## Configuring a queue manager to accept client connections on z/OS:

Before you can run the sample applications, you must first create a queue manager. You can then configure the queue manager to securely accept incoming connection requests from applications that are running in client mode.

### Before you begin

Ensure the queue manager already exists and has been started. Determine whether channel authentication records are already enabled by issuing the MQSC command:

```
DISPLAY QMGR CHLAUTH
```

**Important:** This task expects that channel authentication records are enabled. If this is a queue manager used by other users and applications, changing this setting will affect all other users and applications. If your queue manager does not make use of channel authentication records then step 4 can be replaced with an alternate authentication method (for example a security exit) which sets the MCAUSER to the *non-privileged-user-id* you will obtain in step 1.

You must know which channel name your application expects to use so that the application can be permitted to use the channel. You must also know which objects, for example queues or topics, your application expects to use so that your application can be permitted to use them.

### About this task

This task creates a non-privileged user ID to be used for a client application which connects to the queue manager. Access is granted for the client application only to be able to use the channel it needs and the queue it needs by use of this user ID.

### Procedure

1. Obtain a user ID on the system your queue manager is running on.  
For this task this user ID must not be a privileged administrative user. This user ID is the authority under which the client connection will run on the queue manager.
2. Start a listener program.
  - a. Ensure that your channel initiator is started. If not, start it by issuing the **START CHINIT** command.
  - b. Start the listener program by issuing the following command:  

```
START LISTENER TRPTYPE(TCP) PORT(nnnn)
```

where *nnnn* is your chosen port number.
3. If your application uses the SYSTEM.DEF.SVRCONN then this channel is already defined. If your application uses another channel, create it by issuing the MQSC command:  

```
DEFINE CHANNEL(' channel-name ') CHLTYPE(SVRCONN) TRPTYPE(TCP) +
DESCR('Channel for use by sample programs')
```

*channel-name* is the name of your channel.

4. Create a channel authentication rule allowing only the IP address of your client system to use the channel by issuing the MQSC command:

```
SET CHLAUTH(' channel-name ') TYPE(ADDRESSMAP) ADDRESS(' client-machine-IP-address ') +
MCAUSER(' non-privileged-user-id ')
```

where

*channel-name* is the name of your channel.

*client-machine-IP-address* is the IP address of your client system. If your sample client application is running on the same machine as the queue manager then use an IP address of '127.0.0.1' if your application is going to connect using 'localhost'. If several different client machines are going to connect in, you can use a pattern or a range instead of a single IP address. See Generic IP addresses for details.

*non-privileged-user-id* is the user ID you obtained in step 1 on page 574

5. If your application uses the SYSTEM.DEFAULT.LOCAL.QUEUE, then this queue is already defined. If your application uses another queue, create it by issuing the MQSC command:

```
DEFINE QLOCAL(' queue-name ') DESCR('Queue for use by sample programs')
```

where *queue-name* is the name of your queue.

6. Grant access to connect to and inquire the queue manager:

- a. Ensure that your channel initiator is started. If not, start the channel initiator by issuing the START CHINIT command.

- b. Start a TCP listener, for example issue the following command:

```
START LISTENER TRPTYPE(TCP) PORT(nnnn)
```

where *nnnn* is your chosen port number.

7. If your application is a point-to-point application, that is it makes use of queues, grant access to allow inquiring and the putting and getting messages using your queue by the user ID to be used, by issuing the MQSC commands: Issue the RACF® commands:

```
RDEFINE MQQUEUE qmgr-name.QUEUE. queue-name UACC(NONE)
```

```
PERMIT qmgr-name.QUEUE. queue-name CLASS(MQQUEUE) ID(non-privileged-user-id) ACCESS(UPDATE)
```

where

*qmgr-name* is the name of your queue manager

*queue-name* is the name of your queue.

*non-privileged-user-id* is the user ID you obtained in step 1 on page 574

8. If your application is a publish/subscribe application, that is it makes use of topics, grant access to allow publishing and subscribing using your topic by the user ID to be used, by issuing the following RACF commands:

```
RDEFINE MQTOPIC qmgr-name.PUBLISH.SYSTEM.BASE.TOPIC UACC(NONE)
```

```
PERMIT qmgr-name.PUBLISH.SYSTEM.BASE.TOPIC CLASS(MQTOPIC) ID(non-privileged-user-id) ACCESS(UPDATE)
```

```
RDEFINE MQTOPIC qmgr-name.SUBSCRIBE.SYSTEM.BASE.TOPIC UACC(NONE)
```

```
PERMIT qmgr-name.SUBSCRIBE.SYSTEM.BASE.TOPIC CLASS(MQTOPIC) ID(non-privileged-user-id) ACCESS(UPDATE)
```

where

*qmgr-name* is the name of your queue manager

*non-privileged-user-id* is the user ID you obtained in step 1 on page 574


This will give *non-privileged-user-id* access to any topic in the topic tree, alternatively, you can define a topic object using **DEFINE TOPIC** and grant accesses only to the part of the topic tree referenced by that topic object. For more information, see Controlling user access to topics.

## What to do next

Your client application can now connect to the queue manager and put or get messages using the queue.

### Related information:

SET CHLAUTH  
DEFINE CHANNEL  
DEFINE QLOCAL  
SET AUTHREC

 Authority to work with IBM MQ objects on z/OS

### Preparing and running sample applications for the batch environment on z/OS:

To prepare a sample application that runs in the batch environment, perform the same steps that you would when building any batch IBM MQ for z/OS application.

These steps are listed in “Building z/OS batch applications” on page 421.

Alternatively, where we supply an executable form of a sample, you can run it from the thlqual.SCSQLOAD load library.

**Note:** The assembler language version of the Browse sample uses data control blocks (DCBs), so you must link-edit it using RMODE(24).

The library members to use are listed in Table 85 on page 577, Table 86 on page 577, Table 87 on page 578, and Table 88 on page 578.

You must edit the run JCL supplied for the samples that you want to use (see Table 85 on page 577, Table 86 on page 577, Table 87 on page 578, and Table 88 on page 578 ).

The PARM statement in the supplied JCL contains a number of parameters that you need to modify. To run the C sample programs, separate the parameters by spaces; to run the assembler, COBOL, and PL/I sample programs, separate them by commas. For example, if the name of your queue manager is CSQ1 and you want to run the application with a queue named LOCALQ1, in the COBOL, PL/I, and assembler-language JCL, your PARM statement should look like this:


```
PARM=(CSQ1,LOCALQ1)
```

In the C language JCL, your PARM statement should look like this:

```
PARM=('CSQ1 LOCALQ1')
```

You are now ready to submit the jobs.



Names of the sample batch applications on z/OS: 

A summary of the programs that are supplied for sample batch applications.

The batch application programs are summarized in the following tables:

- Table 85 Put and Get samples
- Table 86 Browse sample
- Table 87 on page 578 Print message sample
- Table 88 on page 578 Publish/Subscribe samples
- Table 89 on page 578 Other samples

*Table 85. Batch Put and Get samples*

| Member name | For language | Description                             | Source file supplied in library | Executable file supplied in library |
|-------------|--------------|-----------------------------------------|---------------------------------|-------------------------------------|
| CSQ4BCJ1    | C            | Get source program                      | SCSQC37S                        | SCSQLOAD                            |
| CSQ4BCK1    | C            | Put source program                      | SCSQC37S                        | SCSQLOAD                            |
| CSQ4BCJR    | C            | Sample run JCL for CSQ4BCJ1 and CSQBCK1 | SCSQPROC                        | None                                |
| CSQ4BVJ1    | COBOL        | Get source program                      | SCSQCOBS                        | SCSQLOAD                            |
| CSQ4BVK1    | COBOL        | Put source program                      | SCSQCOBS                        | SCSQLOAD                            |
| CSQ4BVJR    | COBOL        | Sample run JCL for CSQBVJ1 and CSQBVK1  | SCSQPROC                        | None                                |

*Table 86. Batch Browse sample*

| Member name | For language | Description                 | Source file supplied in library | Executable file supplied in library |
|-------------|--------------|-----------------------------|---------------------------------|-------------------------------------|
| CSQ4BVA1    | COBOL        | Source program              | SCSQCOBS                        | SCSQLOAD                            |
| CSQ4BVAR    | COBOL        | Sample run JCL for CSQ4BVA1 | SCSQPROC                        | None                                |
| CSQ4BAA1    | Assembler    | Source program              | SCSQASMS                        | SCSQLOAD                            |
| CSQ4BAAR    | Assembler    | Sample run JCL for CSQ4BAA1 | SCSQPROC                        | None                                |
| CSQ4BCA1    | C            | Source program              | SCSQC37S                        | SCSQLOAD                            |
| CSQ4BCAR    | C            | Sample run JCL for CSQ4BCA1 | SCSQPROC                        | None                                |
| CSQ4BPA1    | PL/I         | Source program              | SCSQPLIS                        | SCSQLOAD                            |
| CSQ4BPAR    | PL/I         | Sample run JCL for CSQ4BPA1 | SCSQPROC                        | None                                |

Table 87. Batch Print Message sample (C language only)

| Member name | Description                 | Source file supplied in library | Executable file supplied in library |
|-------------|-----------------------------|---------------------------------|-------------------------------------|
| CSQ4BCG1    | Source program              | SCSQC37S                        | SCSQLOAD                            |
| CSQ4BCGR    | Sample run JCL for CSQ4BCG1 | SCSQPROC                        | None                                |
| CSQ4BCL1    | Browse source program       | SCSQC37S                        | SCSQLOAD                            |
| CSQ4BCLR    | Sample run JCL for CSQ4BCL1 | SCSQPROC                        | None                                |

Table 88. Publish/Subscribe samples

| Member name | For language | Description                                                                          | Source file supplied in library | JCL in SCSQPROC | Executable file supplied in library |
|-------------|--------------|--------------------------------------------------------------------------------------|---------------------------------|-----------------|-------------------------------------|
| CSQ4BCP1    | C            | Publish to topic source program                                                      | SCSQC37S                        | CSQ4BCPP        | SCSQLOAD                            |
| CSQ4BCP2    | C            | Subscribe to topic and get messages source program                                   | SCSQC37S                        | CSQ4BCPS        | SCSQLOAD                            |
| CSQ4BCP3    | C            | Subscribe to topic using a user provided destination and get messages source program | SCSQC37S                        | CSQ4BCPD        | SCSQLOAD                            |
| CSQ4BCP4    | C            | Subscribe to topic using extended options and get messages source program            | SCSQC37S                        | CSQ4BCPE        | SCSQLOAD                            |
| CSQ4BVP1    | COBOL        | Publish to topic source program                                                      | SCSQCOBS                        | CSQ4BVPP        | SCSQLOAD                            |
| CSQ4BVP2    | COBOL        | Subscribe to topic and get messages source program                                   | SCSQCOBS                        | CSQ4BVPS        | SCSQLOAD                            |

Table 89. Other samples

| Member name | For language | Description                                       | Source file supplied in library | JCL in SCSQPROC | Executable file supplied in library |
|-------------|--------------|---------------------------------------------------|---------------------------------|-----------------|-------------------------------------|
| CSQ4BCS1    | C            | Asynchronous consumption source program           | SCSQC37S                        | CSQ4BCSC        | SCSQLOAD                            |
| CSQ4BCS2    | C            | Asynchronous Put, and Check status source program | SCSQC37S                        | CSQ4BCSP        | SCSQLOAD                            |
| CSQ4BCM1    | C            | Inquire message properties source program         | SCSQC37S                        | CSQ4BCMP        | SCSQLOAD                            |
| CSQ4BCM2    | C            | Set message properties source program             | SCSQC37S                        | CSQ4BCMP        | SCSQLOAD                            |

## Preparing sample applications for the TSO environment on z/OS:

To prepare a sample application that runs in the TSO environment, perform the same steps that you would when building any batch IBM MQ for z/OS application.

These steps are listed in “Building z/OS batch applications” on page 421. The library members to use are listed in Table 90.

Alternatively, where we supply an executable form of a sample, you can run it from the thlqual.SCSQLOAD load library.

For the Mail Manager sample application, ensure that the queues that it uses are available on your system. They are defined in the member **thlqual.SCSQPROC(CSQ4CVD)**. To ensure that these queues are always available, you could add these members to your CSQINP2 initialization input data set, or use the CSQUTIL program to load these queue definitions.

### Names of the sample TSO applications on z/OS:

Information about the names of the programs that are supplied for each of the sample TSO applications, and the libraries where the source, JCL, and, for the Message Handler sample only, the executable files reside.

The TSO application programs are summarized in the following tables:

- Table 90 Mail manager sample
- Table 91 on page 580 Message handler sample

These samples use ISPF panels. You must therefore include the ISPF stub, ISPLINK, when you link-edit the programs.

*Table 90. TSO Mail Manager sample*

| Member name | For language | Description                          | Source file supplied in library |
|-------------|--------------|--------------------------------------|---------------------------------|
| CSQ4CVD     | independent  | IBM MQ for z/OS object definitions   | SCSQPROC                        |
| CSQ40       | independent  | ISPF messages                        | SCSQMSGE                        |
| CSQ4RVD1    | COBOL        | CLIST to initiate CSQ4TVD1           | SCSQCLST                        |
| CSQ4TVD1    | COBOL        | Source program for Menu program      | SCSQCOBS                        |
| CSQ4TVD2    | COBOL        | Source program for Get Mail program  | SCSQCOBS                        |
| CSQ4TVD4    | COBOL        | Source program for Send Mail program | SCSQCOBS                        |
| CSQ4TVD5    | COBOL        | Source program for Nickname program  | SCSQCOBS                        |
| CSQ4VDP1-6  | COBOL        | Panel definitions                    | SCSQPNLA                        |
| CSQ4VD0     | COBOL        | Data definition                      | SCSQCOBC                        |
| CSQ4VD1     | COBOL        | Data definition                      | SCSQCOBC                        |
| CSQ4VD2     | COBOL        | Data definition                      | SCSQCOBC                        |
| CSQ4VD4     | COBOL        | Data definition                      | SCSQCOBC                        |
| CSQ4RCD1    | C            | CLIST to initiate CSQ4TCD1           | SCSQCLST                        |

Table 90. TSO Mail Manager sample (continued)

| Member name | For language | Description                          | Source file supplied in library |
|-------------|--------------|--------------------------------------|---------------------------------|
| CSQ4TCD1    | C            | Source program for Menu program      | SCSQC37S                        |
| CSQ4TCD2    | C            | Source program for Get Mail program  | SCSQC37S                        |
| CSQ4TCD4    | C            | Source program for Send Mail program | SCSQC37S                        |
| CSQ4TCD5    | C            | Source program for Nickname program  | SCSQC37S                        |
| CSQ4CDP1-6  | C            | Panel definitions                    | SCSQPNLA                        |
| CSQ4TC0     | C            | Include file                         | SCSQC370                        |

Table 91. TSO Message Handler sample

| Member name | For language | Description                            | Source file supplied in library | Executable file supplied in library |
|-------------|--------------|----------------------------------------|---------------------------------|-------------------------------------|
| CSQ4TCH0    | C            | Data definition                        | SCSQC370                        | None                                |
| CSQ4TCH1    | C            | Source program                         | SCSQC37S                        | SCSQLOAD                            |
| CSQ4TCH2    | C            | Source program                         | SCSQC37S                        | SCSQLOAD                            |
| CSQ4TCH3    | C            | Source program                         | SCSQC37S                        | SCSQLOAD                            |
| CSQ4RCH1    | C and COBOL  | CLIST to initiate CSQ4TCH1 or CSQ4TVH1 | SCSQCLST                        | None                                |
| CSQ4CHP1    | C and COBOL  | Panel definition                       | SCSQPNLA                        | None                                |
| CSQ4CHP2    | C and COBOL  | Panel definition                       | SCSQPNLA                        | None                                |
| CSQ4CHP3    | C and COBOL  | Panel definition                       | SCSQPNLA                        | None                                |
| CSQ4CHP9    | C and COBOL  | Panel definition                       | SCSQPNLA                        | None                                |
| CSQ4TVH0    | COBOL        | Data definition                        | SCSQCOBC                        | None                                |
| CSQ4TVH1    | COBOL        | Source program                         | SCSQCOBS                        | SCSQLOAD                            |
| CSQ4TVH2    | COBOL        | Source program                         | SCSQCOBS                        | SCSQLOAD                            |
| CSQ4TVH3    | COBOL        | Source program                         | SCSQCOBS                        | SCSQLOAD                            |

## Preparing the sample applications for the CICS environment on z/OS:

Before you run the CICS sample programs, log on to CICS using a LOGMODE of 32702. This is because the sample programs have been written to use a 3270 mode 2 screen.

To prepare a sample application that runs in the CICS environment, perform the following steps:

1. Create the symbolic description map and the physical screen map for the sample by assembling the BMS screen definition source (supplied in library **thlqual.SCSQMAPS**, where **thlqual** is the high-level qualifier used by your installation). When you name the maps, use the name of the BMS screen definition source (not available for Put and Get sample programs), but omit the last character of that name.
2. Perform the same steps that you would when building any CICS IBM MQ for z/OS application. These steps are listed in “Building CICS applications in z/OS” on page 425. The library members to use are listed in Table 92 on page 582, Table 93 on page 582, Table 94 on page 582, and Table 95 on page 583.

Alternatively, where we supply an executable form of a sample, you can run it from the **thlqual.SCSQCICS** load library.

3. Identify the map set, programs, and transaction to CICS by updating the CICS system definition (CSD) data set. The definitions that you require are in the member **thlqual.SCSQPROC(CSQ4S100)**. For guidance on how to do this, see *The CICS-IBM MQ Adapter* section in the CICS Transaction Server for z/OS Version 4.1 product documentation at: CICS Transaction Server for z/OS Version 4.1, The CICS-IBM MQ adapter.

**Note:** For the Credit Check sample application, you get an error message at this stage if you have not already created the VSAM data set that the sample uses.

4. For the Credit Check and Mail Manager sample applications, ensure that the queues that they use are available on your system. For the Credit Check sample, they are defined in the member **thlqual.SCSQPROC(CSQ4CVB)** for COBOL, and **thlqual.SCSQPROC(CSQ4CCB)** for C. For the Mail Manager sample, they are defined in the member **thlqual.SCSQPROC(CSQ4CVD)**. To ensure that these queues are always available, you could add these members to your CSQINP2 initialization input data set, or use the CSQUTIL program to load these queue definitions.

For the Queue Attributes sample application, you could use one or more of the queues that are supplied for the other sample applications. Alternatively, you could use your own queues. However, in the form that it is supplied, this sample works only with queues that have the characters CSQ4SAMP in the first eight bytes of their name.

## Names of the sample CICS applications on z/OS:

This topic provides a summary of the programs supplied for sample CICS applications.

The CICS application programs are summarized in the following tables:

- Table 92 on page 582 Put and Get samples
- Table 93 on page 582 Queue Attributes sample
- Table 94 on page 582 Mail Manager sample (COBOL only)
- Table 95 on page 583 Credit Check sample
- Table 96 on page 583 Asynchronous Consumption and Publish/Subscribe samples

Table 92. CICS Put and Get samples

| Member name | For language | Description                     | Source file supplied in library | Executable file supplied in library |
|-------------|--------------|---------------------------------|---------------------------------|-------------------------------------|
| CSQ4CCK1    | C            | Put source program              | SCSQC37S                        | SCSQCICS                            |
| CSQ4CCJ1    | C            | Get source program              | SCSQC37S                        | SCSQCICS                            |
| CSQ4CVJ1    | COBOL        | Get source program              | SCSQCOBS                        | SCSQCICS                            |
| CSQ4CVK1    | COBOL        | Put source program              | SCSQCOBS                        | SCSQCICS                            |
| CSQ4S100    | independent  | CICS system definition data set | SCSQPROC                        | None                                |

Table 93. CICS Queue Attributes sample

| Member name | For language | Description                     | Source file supplied in library | Executable file supplied in library |
|-------------|--------------|---------------------------------|---------------------------------|-------------------------------------|
| CSQ4CVC1    | COBOL        | Source program                  | SCSQCOBS                        | SCSQCICS                            |
| CSQ4VMSG    | COBOL        | Message definition              | SCSQCOBC                        | None                                |
| CSQ4VCMS    | COBOL        | BMS screen definition           | SCSQMAPS                        | SCSQCICS (named CSQ4ACM)            |
| CSQ4CAC1    | Assembler    | Source program                  | SCSQASMS                        | SCSQCICS                            |
| CSQ4AMSG    | Assembler    | Message definition              | SCSQMACS                        | None                                |
| CSQ4ACMS    | Assembler    | BMS screen definition           | SCSQMAPS                        | SCSQCICS (named CSQ4ACM)            |
| CSQ4CCC1    | C            | Source program                  | SCSQC37S                        | SCSQCICS                            |
| CSQ4CMMSG   | C            | Message definition              | SCSQC370                        | None                                |
| CSQ4CCMS    | C            | BMS screen definition           | SCSQMAPS                        | SCSQCICS (named CSQ4ACM)            |
| CSQ4S100    | independent  | CICS system definition data set | SCSQPROC                        | None                                |

Table 94. CICS Mail Manager sample (COBOL only)

| Member name | Description                        | Source file supplied in library |
|-------------|------------------------------------|---------------------------------|
| CSQ4CVD     | IBM MQ for z/OS object definitions | SCSQPROC                        |
| CSQ4CVD1    | Source for Menu program            | SCSQCOBS                        |
| CSQ4CVD2    | Source for Get Mail program        | SCSQCOBS                        |
| CSQ4CVD3    | Source for Display Message program | SCSQCOBS                        |
| CSQ4CVD4    | Source for Send Mail program       | SCSQCOBS                        |
| CSQ4CVD5    | Source for Nickname program        | SCSQCOBS                        |
| CSQ4VDMS    | BMS screen definition source       | SCSQMAPS                        |
| CSQ4S100    | CICS system definition data set    | SCSQPROC                        |
| CSQ4VD0     | Data definition                    | SCSQCOBC                        |
| CSQ4VD3     | Data definition                    | SCSQCOBC                        |
| CSQ4VD4     | Data definition                    | SCSQCOBC                        |

Table 95. CICS Credit Check sample

| Member name | For language | Description                                 | Source file supplied in library |
|-------------|--------------|---------------------------------------------|---------------------------------|
| CSQ4CVB     | independent  | IBM MQ object definitions                   | SCSQPROC                        |
| CSQ4CCB     | independent  | IBM MQ object definitions                   | SCSQPROC                        |
| CSQ4CVB1    | COBOL        | Source for user-interface program           | SCSQCOBS                        |
| CSQ4CVB2    | COBOL        | Source for credit application manager       | SCSQCOBS                        |
| CSQ4CVB3    | COBOL        | Source for checking-account program         | SCSQCOBS                        |
| CSQ4CVB4    | COBOL        | Source for distribution program             | SCSQCOBS                        |
| CSQ4CVB5    | COBOL        | Source for agency-query program             | SCSQCOBS                        |
| CSQ4CCB1    | C            | Source for user-interface program           | SCSQ37S                         |
| CSQ4CCB2    | C            | Source for credit application manager       | SCSQ37S                         |
| CSQ4CCB3    | C            | Source for checking-account program         | SCSQ37S                         |
| CSQ4CCB4    | C            | Source for distribution program             | SCSQ37S                         |
| CSQ4CCB5    | C            | Source for agency-query program             | SCSQ37S                         |
| CSQ4CB0     | C            | Include file                                | SCSQ370                         |
| CSQ4CBMS    | C            | BMS screen definition source                | SCSQMAPS                        |
| CSQ4VBMS    | COBOL        | BMS screen definition source                | SCSQMAPS                        |
| CSQ4VB0     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4VB1     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4VB2     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4VB3     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4VB4     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4VB5     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4VB6     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4VB7     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4VB8     | COBOL        | Data definition                             | SCSQCOBC                        |
| CSQ4BAQ     | independent  | Source for VSAM data set                    | SCSQPROC                        |
| CSQ4FILE    | independent  | JCL to build VSAM data set used by CSQ4CVB3 | SCSQPROC                        |
| CSQ4S100    | independent  | CICS system definition data set             | SCSQPROC                        |

Table 96. CICS Asynchronous Consumption and Publish/Subscribe samples

| Member name | Description                                    | Source file supplied in library |
|-------------|------------------------------------------------|---------------------------------|
| CSQ4CVCN    | Source for Simple Message Consumption program  | SCSQCOBS                        |
| CSQ4CVCT    | Source for Control Message Consumption program | SCSQCOBS                        |
| CSQ4CVEV    | Source for Event Handler program               | SCSQCOBS                        |
| CSQ4CVPT    | Source for Message Put Client program          | SCSQCOBS                        |
| CSQ4CVRG    | Source for Registration Client program         | SCSQCOBS                        |
| CSQ4S100    | CICS System Definition data set                | SCSQPROC                        |

## Preparing the sample application for the IMS environment on z/OS: z/OS

Part of the Credit Check sample application can run in the IMS environment.

To prepare this part of the application to run with the CICS sample, first perform the steps described in “Preparing the sample applications for the CICS environment on z/OS” on page 581.

Then perform the following steps:

1. Perform the same steps that you would when building any IMS IBM MQ for z/OS application. These steps are listed in “Building IMS (BMP or MPP) applications” on page 426. The library members to use are listed in Table 97.
2. Identify the application program and database to IMS. Samples are provided with PSBGEN, DBDGEN, ACB definition, IMSGEN, and IMSDALOC statements to enable this.
3. Load the database CSQ4CA by tailoring and running the sample JCL provided for this purpose (CSQ4ILDB). This JCL loads the database with data from the file CSQ4BAQ. Update the IMS control region with a DD statement for the database CSQ4CA.
4. Start the checking-account program as a batch message processing (BMP) program by tailoring and running the sample JCL provided for this purpose. This JCL starts a batch-oriented BMP program. To run the program as a message-oriented BMP program, remove the comment characters from the line in the JCL that contains the IN= statement.

### Names of the sample IMS application on z/OS: z/OS

This information provides a table with the list of the sources and JCLs that are supplied for the Credit Check sample IMS application.

*Table 97. Source and JCL for the Credit Check IMS sample (C only)*

| Member name | Description                                             | Supplied in library |
|-------------|---------------------------------------------------------|---------------------|
| CSQ4CVB     | IBM MQ object definitions                               | SCSQPROC            |
| CSQ4ICB3    | Source for checking-account program                     | SCSQC37S            |
| CSQ4ICBL    | Source for loading the checking-account database        | SCSQC37S            |
| CSQ4CBI     | Data definition                                         | SCSQC370            |
| CSQ4PSBL    | PSBGEN JCL for database-load program                    | SCSQPROC            |
| CSQ4PSB3    | PSBGEN JCL for checking-account program                 | SCSQPROC            |
| CSQ4DBDS    | DBDGEN JCL for database CSQ4CA                          | SCSQPROC            |
| CSQ4GIMS    | IMSGEN macro definitions for CSQ4IVB3 and CSQ4CA        | SCSQPROC            |
| CSQ4ACBG    | Application control block (ACB) definition for CSQ4IVB3 | SCSQPROC            |
| CSQ4BAQ     | Source for database                                     | SCSQPROC            |
| CSQ4ILDB    | Sample run JCL for database-load job                    | SCSQPROC            |
| CSQ4ICBR    | Sample run JCL for checking-account program             | SCSQPROC            |
| CSQ4DYNA    | IMSDALOC macro definitions for database                 | SCSQPROC            |



## The Put samples on z/OS:

The Put sample programs put messages on a queue using the MQPUT call.

The source programs are supplied in C and COBOL in the batch and CICS environments (see Table 85 on page 577 and Table 92 on page 582 ).

### Design of the Put sample

The flow through the program logic is:

1. Connect to the queue manager using the MQCONN call. If this call fails, print the completion and reason codes and stop processing.

**Note:** If you are running the sample in a CICS environment, you do not need to issue an MQCONN call; if you do, it returns DEF\_HCONN. You can use the connection handle MQHC\_DEF\_HCONN for the MQI calls that follow.

2. Open the queue using the MQOPEN call with the MQOO\_OUTPUT option. On input to this call, the program uses the connection handle that is returned in step 1 on page 587. For the object descriptor structure (MQOD), it uses the default values for all fields except the queue name field, which is passed as a parameter to the program. If the MQOPEN call fails, print the completion and reason codes and stop processing.
3. Create a loop within the program issuing MQPUT calls until the required number of messages are put on the queue. If an MQPUT call fails, the loop is abandoned early, no further MQPUT calls are attempted, and the completion and reason codes are returned.
4. Close the queue using the MQCLOSE call with the object handle returned in step 2 on page 587. If this call fails, print the completion and reason codes.
5. Disconnect from the queue manager using the MQDISC call with the connection handle returned in step 1 on page 587. If this call fails, print the completion and reason codes.

**Note:** If you are running the sample in a CICS environment, you do not need to issue an MQDISC call.

### The Put samples for the batch environment on z/OS:

Use this topic when considering Put samples for the batch environment.

To run the samples, edit and run the sample JCL, as described in “Preparing and running sample applications for the batch environment on z/OS” on page 576.

The programs take the following parameters in an EXEC PARM, separated by spaces in C and commas in COBOL:


1. The name of the queue manager (4 characters)
2. The name of the target queue (48 characters)
3. The number of messages (up to 4 digits)
4. The padding character to write in the message (1 character)
5. The number of characters to write in the message (up to 4 digits)
6. The persistence of the message (1 character: P for persistent or N for nonpersistent)

If you enter any of the these parameters wrongly, you receive appropriate error messages.

Any messages from the samples are written to the SYSPRINT data set.

## Usage notes

- To keep the samples simple, there are some minor functional differences between language versions. However, these differences are minimized if you use the layout of the parameters shown in the sample run JCL, CSQ4BCJR, and CSQ4BVJR. None of the differences relate to the MQI.
- CSQ4BCK1 allows you to enter more than four digits for the number of messages sent and the length of the messages.
- For the two numeric fields, enter any digit in the range 1 through 9999. The value that you enter should be a positive number. For example, to put a single message, you can enter 1, 01, 001, or 0001 as the value. If you enter nonnumeric or negative values, you might receive an error. For example, if you enter -1, the COBOL program sends a 1-byte message, but the C program receives an error.
- For both programs, CSQ4BCK1 and CSQ4BVK1, you must enter P in the persistence parameter, ++PER++, if you want the message to be persistent. If you fail to do so, the message will be nonpersistent.

The Put samples for the CICS environment on z/OS: 

Use this topic when considering Put samples for the CICS environment.

The transactions take the following parameters separated by commas:

1. The number of messages (up to 4 digits)
2. The padding character to write in the message (1 character)
3. The number of characters to write in the message (up to 4 digits)
4. The persistence of the message (1 character: P for persistent or N for nonpersistent)
5. The name of the target queue (48 characters)

If you enter any of these parameters wrongly, you receive appropriate error messages.

For the COBOL sample, invoke the Put sample in the CICS environment by entering:

```
MVPT,9999,*,9999,P,QUEUE.NAME
```

For the C sample, invoke the Put sample in the CICS environment by entering:

```
MCPT,9999,*,9999,P,QUEUE.NAME
```

Any messages from the samples are displayed on the screen.

## Usage notes

- To keep the samples simple, there are some minor functional differences between language versions. None of the differences relate to the MQI.
- If you enter a queue name that is longer than 48 characters, its length is truncated to the maximum of 48 characters but no error message is returned.
- Before entering the transaction, press the CLEAR key.
- For the two numeric fields, enter any number in the range 1 through 9999. The value that you enter should be a positive number. For example, to put a single message, you can enter the value 1, 01, 001, or 0001. If you enter nonnumeric or negative values, you might receive an error. For example, if you enter -1, the COBOL program sends a 1-byte message, and the C program abends with an error from malloc().
- For both programs, CSQ4CCK1 and CSQ4CVK1, enter P in the persistence parameter if you want the message to be persistent. For non-persistent messages, enter N in the persistence parameter. If you enter any other value you receive an error message.
- The messages are put in syncpoint because default values are used for all parameters except those set during program invocation.

## The Get samples on z/OS:

The Get sample programs get messages from a queue using the MQGET call.

The source programs are supplied in C and COBOL in the batch and CICS environments (see Table 85 on page 577 and Table 92 on page 582 ).

## Design of the Get sample on z/OS:

Learn about the design of the Get sample, and some usage notes to consider.

The flow through the program logic is:

1. Connect to the queue manager using the MQCONN call. If this call fails, print the completion and reason codes and stop processing.

**Note:** If you are running the sample in a CICS environment, you do not need to issue an MQCONN call; if you do, it returns DEF\_HCONN. You can use the connection handle MQHC\_DEF\_HCONN for the MQI calls that follow.

2. Open the queue using the MQOPEN call with the MQOO\_INPUT\_SHARED and MQOO\_BROWSE options. On input to this call, the program uses the connection handle that is returned in step 1. For the object descriptor structure (MQOD), it uses the default values for all fields except the queue name field, which is passed as a parameter to the program. If the MQOPEN call fails, print the completion and reason codes and stop processing.
3. Create a loop within the program issuing MQGET calls until the required number of messages are retrieved from the queue. If an MQGET call fails, the loop is abandoned early, no further MQGET calls are attempted, and the completion and reason codes are returned. The following options are specified on the MQGET call:
  - MQGMO\_NO\_WAIT
  - MQGMO\_ACCEPT\_TRUNCATED\_MESSAGE
  - MQGMO\_SYNCPOINT or MQGMO\_NO\_SYNCPOINT
  - MQGMO\_BROWSE\_FIRST and MQGMO\_BROWSE\_NEXT

For a description of these options, see MQGET. For each message, the message number is printed followed by the length of the message and the message data.

4. Close the queue using the MQCLOSE call with the object handle returned in step 2. If this call fails, print the completion and reason codes.
5. Disconnect from the queue manager using the MQDISC call with the connection handle returned in step 1. If this call fails, print the completion and reason codes.


**Note:** If you are running the sample in a CICS environment, you do not need to issue an MQDISC call.

## Usage notes

- To keep the samples simple, there are some minor functional differences between language versions. However, these differences are minimized if you use the layout of the parameters shown in the sample run JCL, CSQ4BCJR, and CSQ4BVJR. None of the differences relate to the MQI.
- CSQ4BCJ1 allows you to enter more than four digits for the number of messages retrieved.
- Messages longer than 64 KB are truncated.
- CSQ4BCJ1 can only correctly display character messages because it only displays until the first NULL (\0) character is displayed.
- For the numeric number-of-messages field, enter any digit in the range 1 through 9999. The value that you enter should be a positive number. For example, to get a single message, you can enter 1, 01, 001,

or 0001 as the value. If you enter nonnumeric or negative values, you might receive an error. For example, if you enter -1, the COBOL program retrieves one message, but the C program does not retrieve any messages.

- For both programs, CSQ4BCJ1 and CSQ4BVJ1, enter B in the get parameter, ++GET++, if you want to browse the messages.
- For both programs, CSQ4BCJ1 and CSQ4BVJ1, enter S in the syncpoint parameter, ++SYNC++, for messages to be retrieved in syncpoint.

The Get samples for the batch environment on z/OS: 

To run the samples, edit and run the sample JCL, as described in “Preparing and running sample applications for the batch environment on z/OS” on page 576.


The programs take the following parameters in an EXEC PARM, separated by spaces in C and commas in COBOL:

1. The name of the queue manager (4 characters)
2. The name of the target queue (48 characters)
3. The number of messages to get (up to 4 digits)
4. The browse/get message option (1 character: B to browse or D to destructively get the messages)
5. The syncpoint control (1 character: S for syncpoint or N for no syncpoint)

If you enter any of these parameters incorrectly, you receive appropriate error messages.

Output from the samples is written to the SYSPRINT data set:

```
=====
PARAMETERS PASSED :
QMGR - VC9
QNAME - A.Q
NUMMSGs - 000000002
GET - D
SYNCPOINT - N
=====
MQCONN SUCCESSFUL
MQOPEN SUCCESSFUL
000000000 : 000000010 : *****
000000001 : 000000010 : *****
000000002 MESSAGES GOT FROM QUEUE
MQCLOSE SUCCESSFUL
MQDISC SUCCESSFUL
```

The Get samples for the CICS environment on z/OS: 

Special considerations for the Get samples for the CICS environment.

The transactions take the following parameters in an EXEC PARM, separated by commas:

1. The number of messages to get (up to four digits)
2. The browse/get message option (one character: B to browse or D to destructively get the messages)
3. The syncpoint control (one character: S for syncpoint or N for no syncpoint)
4. The name of the target queue (48 characters)

If you enter any of these parameters incorrectly, you receive appropriate error messages.

For the COBOL sample, invoke the Get sample in the CICS environment by entering:

```
MVGT,9999,B,S,QUEUE.NAME
```

For the C sample, invoke the Get sample in the CICS environment by entering:

```
MCGT,9999,B,S,QUEUE.NAME
```

When the messages are retrieved from the queue, they are put on a CICS temporary storage queue with the same name as the CICS transaction (for example, MCGT for the C sample).

Here is example output of the Get samples:

```
***** TOP OF QUEUE *****
00000000 : 00000010: *****
00000001 : 00000010 :*****
***** BOTTOM OF QUEUE *****
```

### Usage notes

- To keep the samples simple, there are some minor functional differences between language versions. None of the differences relate to the MQI.
- If you enter a queue name that is longer than 48 characters, its length is truncated to the maximum of 48 characters but no error message is returned.
- Before entering the transaction, press the CLEAR key.
- CSQ4CCJ1 can only correctly display character messages because it only displays until the first NULL (\0) character is displayed.
- For the numeric field, enter any number in the range 1 through 9999. The value that you enter should be a positive number. For example, to get a single message, you can enter the value 1, 01, 001, or 0001. If you enter a nonnumeric or negative value, you might receive an error.
- Messages longer than 24 526 bytes in C and 9 950 bytes in COBOL are truncated. This is due to the way that the CICS temporary storage queues are used.
- For both programs, CSQ4CCK1 and CSQ4CVK1, enter B in the get parameter if you want to browse the messages, otherwise enter D. This performs destructive MQGET calls. If you enter any other value you receive an error message.
- For both programs, CSQ4CCJ1 and CSQ4CVJ1, enter S in the syncpoint parameter to retrieve messages in syncpoint. If you enter N in the syncpoint parameter, the MQGET calls are issued out of syncpoint. If you enter any other value you receive an error message.

## The Browse sample on z/OS:

The Browse sample is a batch application that demonstrates how to browse messages on a queue using the MQGET call.

The application steps through all the messages in a queue, printing the first 80 bytes of each one. You could use this application to look at the messages on a queue without changing them.

Source programs and sample run JCL are supplied in the COBOL, assembler, PL/I, and C languages (see Table 86 on page 577 ).

To start the application, edit and run the sample run JCL, as described in “Preparing and running sample applications for the batch environment on z/OS” on page 576. You can look at messages on one of your own queues by specifying the name of the queue in the run JCL.

When you run the application (and there are some messages on the queue), the output data set looks this:

```
07/12/1998 SAMPLE QUEUE REPORT PAGE 1
QUEUE MANAGER NAME : VC4
QUEUE NAME : CSQ4SAMP.DEAD.QUEUE
RELATIVE
MESSAGE MESSAGE
NUMBER LENGTH ----- MESSAGE DATA -----
1 740 HELLO. PLEASE CALL ME WHEN YOU GET BACK.
2 429 CSQ4BQRM
3 429 CSQ4BQRM
4 429 CSQ4BQRM
5 22 THIS IS A TEST MESSAGE
6 8 CSQ4TEST
7 36 CSQ4MSG - ANOTHER TEST MESSAGE.....
!8 9 CSQ4STOP
***** END OF REPORT *****
```

If there are no messages on the queue, the data set contains the headings and the End of report message only. If an error occurs with any of the MQI calls, the completion and reason codes are added to the output data set.

## Design of the Browse sample on z/OS:

The Browse sample application uses a single program module; one is provided in each of the supported programming languages.

The flow through the program logic is:

1. Open a print data set and print the title line of the report. Check that the names of the queue manager and queue have been passed from the run JCL. If both names have been passed, print the lines of the report that contain the names. If they have not, print an error message, close the print data set, and stop processing.

The way that the program tests the parameters it is passed from the JCL depends on the language in which the program is written; for more information, see “Language-dependent design considerations on z/OS” on page 592.

2. Connect to the queue manager using the MQCONN call. If this call is not successful, print the completion and reason codes, close the print data set, and stop processing.
3. Open the queue using the MQOPEN call with the MQOO\_BROWSE option. On input to this call, the program uses the connection handle returned in step 2. For the object descriptor structure (MQOD), it

uses the default values for all the fields except the queue name (which was passed in step 1 on page 590 ). If this call is not successful, print the completion and reason codes, close the print data set, and stop processing.

4. Browse the first message on the queue, using the MQGET call. On input to this call, the program specifies:

- The connection and queue handles from steps 2 on page 590 and 3 on page 590
- An MQMD structure with all fields set to their initial values
- Two options:
  - MQGMO\_BROWSE\_FIRST
  - MQGMO\_ACCEPT\_TRUNCATED\_MSG
- A buffer of size 80 bytes to hold the data copied from the message

The MQGMO\_ACCEPT\_TRUNCATED\_MSG option allows the call to complete even if the message is longer than the 80-byte buffer specified in the call. If the message is longer than the buffer, the message is truncated to fit the buffer, and the completion and reason codes are set to show this. The sample was designed so that messages are truncated to 80 characters to make the report easy to read. The buffer size is set by a DEFINE statement, so you can easily change it if you want to.

5. Perform the following loop until the MQGET call fails:

- a. Print a line of the report showing:

- The sequence number of the message (this is a count of the browse operations).
- The true length of the message (not the truncated length). This value is returned in the DataLength field of the MQGET call.
- The first 80 bytes of the message data.

- b. Reset the MsqId and CorrelId fields of the MQMD structure to nulls

- c. Browse the next message, using the MQGET call with these two options:

- MQGMO\_BROWSE\_NEXT
- MQGMO\_ACCEPT\_TRUNCATED\_MSG

6. If the MQGET call fails, test the reason code to see if the call has failed because the browse cursor has got to the end of the queue. In this case, print the End of report message and go to step 7 ; otherwise, print the completion and reason codes, close the print data set, and stop processing.
7. Close the queue using the MQCLOSE call with the object handle returned in step 3 on page 590.
8. Disconnect from the queue manager using the MQDISC call with the connection handle returned in step 2 on page 590.
9. Close the print data set and stop processing.

Source modules are provided for the Browse sample in four programming languages.

There are two main differences between the source modules:

- When testing the parameters passed from the run JCL, the COBOL, PL/I, and assembler-language modules search for the comma character (.). If the JCL passes PARM=(,LOCALQ1), the application attempts to open queue LOCALQ1 on the default queue manager. If there is no name after the comma (or no comma), the application returns an error. The C module does not search for the comma character. If the JCL passes a single parameter (for example, PARM=('LOCALQ1')), the C module uses this as a queue name on the default queue manager.
- To keep the assembler-language module simple, it uses the date format *yy/ddd* (for example, 05/116) when it creates the print report. The other modules use the calendar date in *mm/dd/yy* format.

### The Print Message sample on z/OS: z/OS

The Print Message sample is a batch application that demonstrates how to remove all the messages from a queue using the MQGET call.

The Print Message sample uses three parameters:

1. The name of the queue manager
2. The name of the source queue
3. An optional parameter for properties

It also prints, for each message, the fields of the message descriptor, followed by the message data. The program prints the data both in hexadecimal and as characters (if they are printable). If a character is not printable, the program replaces it with a period character (.). You can use the program when diagnosing problems with an application that is putting messages on a queue.

Permissible values for the property parameter are:

*Table 98. Permissible values for the property parameter*

| Value | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 0     | Default behavior, as it was for V6. The properties that get delivered to the application depend on the <b>PropertyControl</b> queue attribute that the message is retrieved from.                                                                                                                                                                                                                                                                                                                                                                                                            |
| 1     | A message handle is created and used with the MQGET. Properties of the message, except those contained in the message descriptor (or extension) are displayed in a similar fashion to the message descriptor. For example:<br><pre>****Message properties****   property name: property value</pre> Or if no properties are available:<br><pre>****Message properties****   None</pre> Numeric values are displayed using printf, string values are surrounding in single quotation marks, and byte strings are surrounded with X and single quotation marks, as for the message descriptor. |
| 2     | MQGMO_NO_PROPERTIES is specified, so that only message descriptor properties will be returned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 3     | MQGMO_PROPERTIES_FORCE_MQRFH2 is specified, so that all properties are returned in the message data.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 4     | MQGMO_PROPERTIES_COMPATIBILITY is specified, so that all properties can be returned depending on whether a Version 6 property is included, otherwise the properties are discarded.                                                                                                                                                                                                                                                                                                                                                                                                           |



You can change the application so that it browses the messages, rather than removing them from the queue. To do this, compile with the option of `-DBROWSE`, to define the `BROWSE` macro, as indicated in “Design of the Print Message sample on z/OS” on page 594. Executable code is provided for you in the `SCSQLOAD` library. Module `CSQ4BCG0` is built with `-DBROWSE`; module `CSQ4BCG1` destructively reads the queue.

The application has a single source program, which is written in the C language. Sample run JCL code is also supplied (see Table 87 on page 578 ).

To start the application, edit and run the sample run JCL, as described in “Preparing and running sample applications for the batch environment on z/OS” on page 576. When you run the application (and there are some messages on the queue), the output data set looks like that in Figure 94 on page 594.



3. Open the queue using the MQOPEN call with the MQOO\_INPUT\_SHARED option.

**Note:** If you want the application to browse the messages rather than remove them from the queue, compile the sample with -DBROWSE, or, add #define BROWSE at the beginning of the source. When you do this, the macro preprocessor adds the line in the program that selects the MQOO\_BROWSE option in the compilation.

On input to this call, the program uses the connection handle returned in step 2 on page 594. For the object descriptor structure (MQOD), it uses the default values for all the fields except the queue name (which was passed in step 1 on page 594 ). If this call is not successful, print the completion and reason codes and stop processing; otherwise, print the name of the queue.

4. If you use a message handle to obtain the message properties use MQCRTMH to create such a handle for use with subsequent MQGET calls. If this call is not successful, print the completion and reason codes and stop processing.
5. Set the get message options to reflect the request action for any message properties.
6. Perform the following loop until the MQGET call fails:
  - a. Initialize the buffer to blanks so that the message data does not get corrupted by any data already in the buffer.
  - b. Set the MsgId and CorrelId fields of the MQMD structure to nulls so that the MQGET call selects the first message from the queue.
  - c. Get a message from the queue, using the MQGET call. On input to this call, the program specifies:
    - The connection and object handles from steps 2 on page 594 and 3.
    - An MQMD structure with all fields set to their initial values. (MsgId and CorrelId are reset to nulls for each MQGET call.)
    - The option MQGMO\_NO\_WAIT.

**Note:** If you want the application to browse the messages rather than remove them from the queue, compile the sample with -DBROWSE, or, add #define BROWSE at the beginning of the source. When you do this, the macro preprocessor adds the line in the program that selects the MQGMO\_BROWSE\_NEXT option to the compilation. When this option is used on a call against a queue for which no browse cursor has previously been used with the current object handle, the browse cursor is positioned logically before the first message.

- d. Call the printMD subroutine. This prints the name of each field in the message descriptor, followed by its contents.
  - e. If you created a message handle in step 4 call the printProperties subroutine to display any message properties.
  - f. Print the length of the message, followed by the message data. Each line of message data is in this format:
    - Relative position (in hexadecimal) of this part of the data
    - 16 bytes of hexadecimal data
    - The same 16 bytes of data in character format, if it is printable (nonprintable characters are replaced by periods)
7. If the MQGET call fails, test the reason code to see if the call failed because there are no more messages on the queue. In this case, print the message: No more messages; otherwise, print the completion and reason codes. In both cases, go to step 9 on page 596.

**Note:** The MQGET call fails if it finds a message that has more than 64KB of data. To change the program to handle larger messages, you could do one of the following:

- Add the MQGMO\_ACCEPT\_TRUNCATED\_MSG option to the MQGET call, so that the call gets the first 64KB of data and discards the remainder

- Make the program leave the message on the queue when it finds one with this amount of data
  - Increase the size of the buffer
8. If you created a message handle in step 4 on page 595 call MQDLTMH to delete it.
  9. Close the queue using the MQCLOSE call with the object handle returned in step 3 on page 595.
  10. Disconnect from the queue manager using the MQDISC call with the connection handle returned in step 2 on page 594.

### The Queue Attributes sample on z/OS:

The Queue Attributes sample is a conversational-mode CICS application that demonstrates the use of the MQINQ and MQSET calls.

It shows how to inquire about the values of the **InhibitPut** and **InhibitGet** attributes of queues, and how to change them so that programs cannot put messages on, or get messages from, a queue. You might want to *lock* a queue in this way when you are testing a program.

To prevent accidental interference with your own queues, this sample works only on a queue object that has the characters CSQ4SAMP in the first eight bytes of its name. However, the source code includes comments to show you how to remove this restriction.

Source programs are supplied in the COBOL, assembler, and C languages (see Table 93 on page 582 ).

The assembler-language version of the sample uses reenterable code. To do this, you will notice that the code for each MQI call in that version of the sample includes the MF keyword; for example:

```
CALL MQCONN, (NAME, HCONN, COMPCODE, REASON), MF=(E, PARMAREA), VL
```

(The VL keyword means that you can use the CICS Execution Diagnostic Facility (CEDF) supplied transaction for debugging the program.) For more information about writing reenterable programs, see Coding in System/390 assembler language.

To start the application, start your CICS system and use the following CICS transactions:

- For COBOL, MVC1
- For assembler language, MAC1
- For C, MCC1

You can change the name of any of these transactions by changing the CSD data set mentioned in step 3.

### Design of the sample

When you start the sample, it displays a screen map that has fields for:

- Name of the queue
- User request (valid actions are: inquire, allow, or inhibit)
- Current status of put operations for the queue
- Current status of get operations for the queue

The first two fields are for user input. The last two fields are filled by the application: they show the word INHIBITED or the word ALLOWED.

The application validates the values that you enter in the first two fields. It checks that the queue name starts with the characters CSQ4SAMP and that you entered one of the three valid requests in the Action field. The application converts all your input to uppercase, so you cannot use any queues with names that contain lowercase characters.

If you enter inquire in the **Action** field, the flow through the program logic is:

1. Open the queue using the MQOPEN call with the MQOO\_INQUIRE option
2. Call MQINQ using the selectors MQIA\_INHIBIT\_GET and MQIA\_INHIBIT\_PUT
3. Close the queue using the MQCLOSE call
4. Analyze the attributes that are returned in the **IntAttrs** parameter of the MQINQ call and move the words INHIBITED or ALLOWED, as appropriate, to the relevant screen fields

If you enter `inhibit` in the **Action** field, the flow through the program logic is:

1. Open the queue using the MQOPEN call with the MQOO\_SET option
2. Call MQSET using the selectors MQIA\_INHIBIT\_GET and MQIA\_INHIBIT\_PUT, and with the values MQQA\_GET\_INHIBITED and MQQA\_PUT\_INHIBITED in the **IntAttrs** parameter
3. Close the queue using the MQCLOSE call
4. Move the word INHIBITED to the relevant screen fields

If you enter `allow` in the **Action** field, the application performs similar processing to that for an inhibit request. The only differences are the settings of the attributes and the words displayed on the screen.

When the application opens the queue, it uses the default connection handle to the queue manager. (CICS establishes a connection to the queue manager when you start your CICS system.) The application can trap the following errors at this stage:

- The application is not connected to the queue manager
- The queue does not exist
- The user is not authorized to access the queue
- The application is not authorized to open the queue

For other MQI errors, the application displays the completion and reason codes.


#### The Mail Manager sample on z/OS:

The Mail Manager sample application is a suite of programs that demonstrates sending and receiving messages, both within a single environment and across different environments. The application is a simple electronic mailing system that allows users to exchange messages, even if they use different queue managers.

The application demonstrates how to create queues using the MQOPEN call and by putting IBM MQ for z/OS commands on the system-command input queue.

Three versions of the application are provided:

- A CICS application written in COBOL
- A TSO application written in COBOL
- A TSO application written in C

*Preparing the Mail Manager sample on z/OS:* 

The Mail Manager is provided in versions that run in two environments. The preparation that you must carry out before you run the application depends on the environment that you want to use.

Users can access mail queues and nickname queues from both TSO and CICS so long as their sign-on user IDs are the same on each system.

Before you can send messages to another queue manager, you must set up a message channel to that queue manager. To do this, use the channel control function of IBM MQ, described in Channel control function.

### **Preparing the sample for the TSO environment**

Follow these steps:

1. Prepare the sample as described in “Preparing sample applications for the TSO environment on z/OS” on page 579.

2. Tailor the CLIST provided for the sample to define:


- The location of the panels
- The location of the message file
- The location of the load modules
- The name of the queue manager that you want to use with the application

A separate CLIST is provided for each language version of the sample:

- For the COBOL version: CSQ4RVD1
- For the C version: CSQ4RCD1

3. Ensure that the queues used by the application are available on the queue manager. (The queues are defined in CSQ4CVD.)

**Note:** VS COBOL II does not support multitasking with ISPF. This means that you cannot use the Mail Manager sample application on both sides of a split screen. If you do, the results are unpredictable.

*Running the Mail Manager sample on z/OS:* 

To start the sample in the CICS Transaction Server for z/OS environment, run transaction MAIL. If you have not already signed on to CICS, the application prompts you to enter a user ID to which it can send your mail.

When you start the application, it opens your mail queue. If this queue does not exist, the application creates one for you. Mail queues have names of the form CSQ4SAMP.MAILMGR. *userid*, where *userid* depends on the environment:

#### **In TSO**

The user's TSO ID

#### **In CICS**

The user's CICS sign-on or the user ID entered by the user when prompted when the Mail Manager started

All parts of the queue names that the Mail Manager uses must be uppercase.

The application then presents a menu panel that has options for:

- Read incoming mail
- Send mail
- Create nickname

The menu panel also shows you how many messages are waiting on your mail queue. Each of the menu options displays a further panel:

### Read incoming mail

The Mail Manager displays a list of the messages that are on your mail queue. (Only the first 99 messages on the queue are displayed.) For an example of this panel, see Figure 97 on page 602. When you select a message from this list, the contents of the message are displayed (see Figure 98 on page 603 ).

### Send mail

A panel prompts you to enter:

- The name of the user to whom you want to send a message
- The name of the queue manager that owns their mail queue
- The text of your message

In the user name field, you can enter either a user ID or a nickname that you created using the Mail Manager. You can leave the queue manager name field blank if the user's mail queue is owned by the same queue manager that you are using, and you must leave it blank if you entered a nickname in the user name field:

- If you specify only a user name, the program first assumes that the name is a nickname, and sends the message to the object defined by that name. If there is no such nickname, the program attempts to send the message to a local queue of that name.
- If you specify both a user name and a queue manager name, the program sends the message to the mail queue that is defined by those two names.

For example, if you want to send a message to user JONESM on remote queue manager QM12, you could send them a message in either of two ways:

- Use both fields to specify user JONESM at queue manager QM12.
- Define a nickname (for example, MARY) for that user and send them a message by putting MARY in the user name field and nothing in the queue manager name field.

### Create nickname

You can define an easy-to-remember name that you can use when you send a message to another user who you contact frequently. You are prompted to enter the user ID of the other user and the name of the queue manager that owns their mail queue.

Nicknames are queues that have names of the form `CSQ4SAMP.MAILMGR.userid.nickname`, where *userid* is your own user ID and *nickname* is the nickname that you want to use. With names structured in this way, users can each have their own set of nicknames.

The type of queue that the program creates depends on how you complete the fields of the Create Nickname panel:

- If you specify only a user name, or the queue manager name is the same as that of the queue manager to which the Mail Manager is connected, the program creates an alias queue.
- If you specify both a user name and a queue manager name (and the queue manager is not the one to which the Mail Manager is connected), the program creates a local definition of a remote queue. The program does not check the existence of the queue to which this definition resolves, or even that the remote queue manager exists.

For example, if your own user ID is SMITHK and you create a nickname called MARY for user JONESM (who uses the remote queue manager QM12), the nickname program creates a local definition of a remote queue named `CSQ4SAMP.MAILMGR.SMITHK.MARY`. This definition resolves to Mary's mail queue, which is `CSQ4SAMP.MAILMGR.JONESM` at queue manager QM12. If you are using queue manager QM12 yourself, the program instead creates an alias queue of the same name (`CSQ4SAMP.MAILMGR.SMITHK.MARY`).

The C version of the TSO application makes greater use of ISPF's message-handling capabilities than does the COBOL version. You might notice that different error messages are displayed by the C and COBOL

versions.

Design of the Mail Manager sample on z/OS: z/OS

The following sections describe each of the programs that make up the Mail Manager sample application.

The relationships between the programs and the panels that the application uses is shown in Figure 95 for the TSO version, and Figure 96 on page 601 for the CICS Transaction Server for z/OS version.

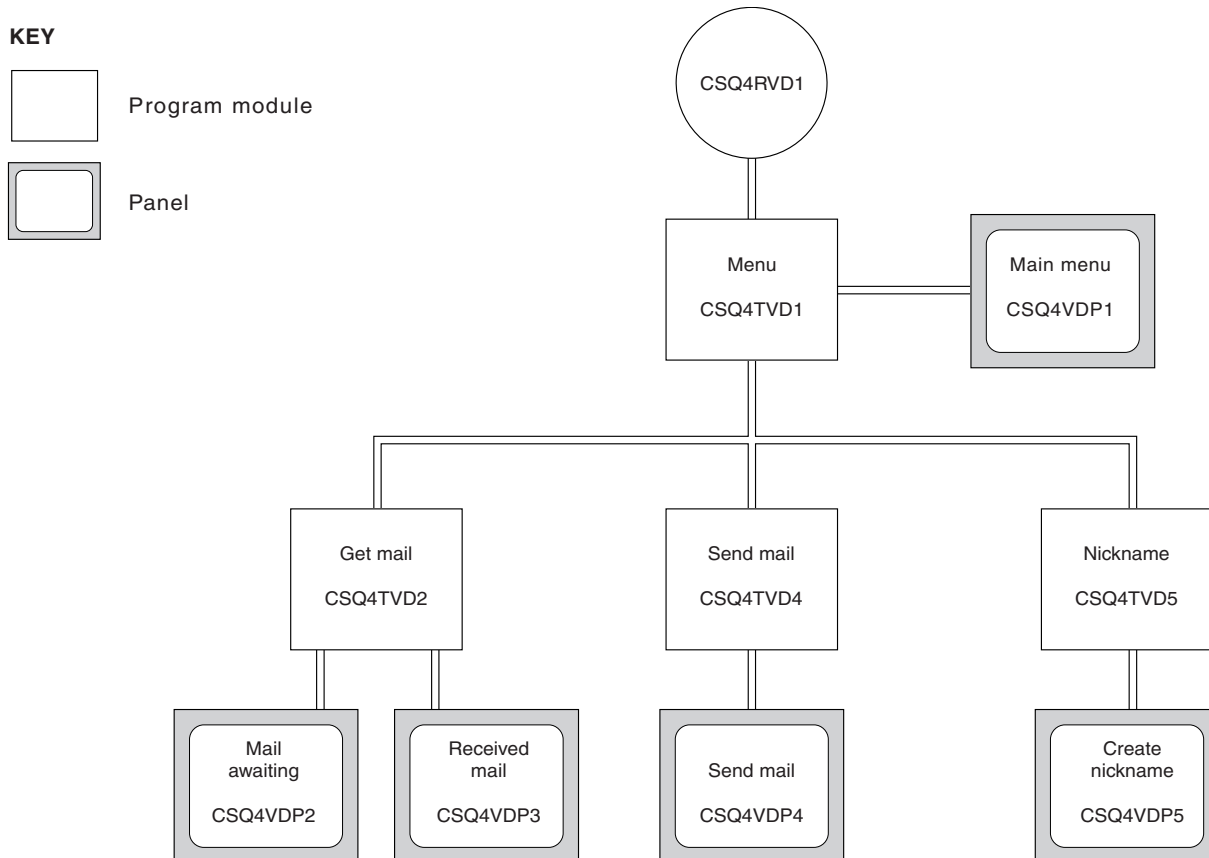


Figure 95. Programs and panels for the TSO versions of the Mail Manager. This figure shows the names for the COBOL version.



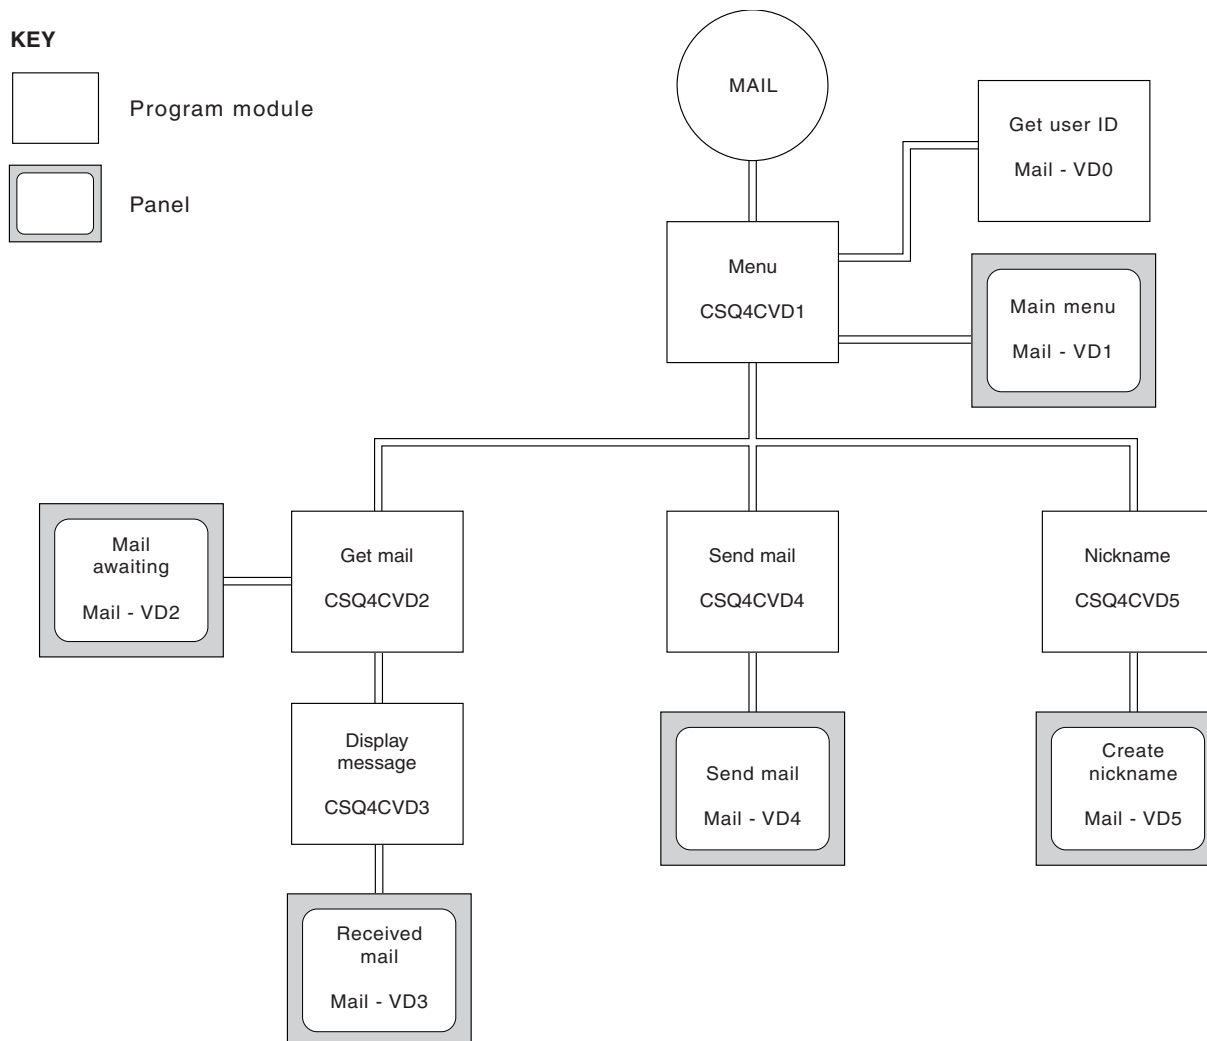



Figure 96. Programs and panels for the CICS version of the Mail Manager

Menu program on z/OS: 

In the TSO environment, the menu program is invoked by the CLIST. In the CICS environment, the program is invoked by transaction MAIL.

The menu program (CSQ4TVD1 for TSO, CSQ4CVD1 for CICS) is the initial program in the suite. It displays the menu (CSQ4VDP1 for TSO, VD1 for CICS) and invokes the other programs when they are selected from the menu.

The program first obtains the user's ID:

- In the CICS version of the program, if the user has signed on to CICS, the user ID is obtained by using the CICS command ASSIGN USERID. If the user has not signed on, the program displays the sign on panel (CSQ4VD0) to prompt the user to enter a user ID. There is no security processing within this program; the user can give any user ID.
- In the TSO version, the user's ID is obtained from TSO in the CLIST. It is passed to the menu program as a variable in the ISPF shared pool.

After the program has obtained the user ID, it checks to ensure that the user has a mail queue (CSQ4SAMP.MAILMGR. *userid* ). If a mail queue does not exist, the program creates one by putting a

message on the system-command input queue. The message contains the IBM MQ for z/OS command DEFINE QLOCAL. The object definition that this command uses sets the maximum depth of the queue to 9999 messages.

The program also creates a temporary dynamic queue to handle replies from the system-command input queue. To do this, the program uses the MQOPEN call, specifying the SYSTEM.DEFAULT.MODEL.QUEUE as the template for the dynamic queue. The queue manager creates the temporary dynamic queue with a name that has the prefix CSQ4SAMP; the remainder of the name is generated by the queue manager.

The program then opens the user's mail queue and finds the number of messages on the queue by inquiring about the current depth of the queue. To do this, the program uses the MQINQ call, specifying the MQIA\_CURRENT\_Q\_DEPTH selector.

The program then performs a loop that displays the menu and processes the selection that the user makes. The loop is stopped when the user presses the PF3 key. When a valid selection is made, the appropriate program is started; otherwise an error message is displayed.

Get-mail and display-message programs on z/OS: 

In the TSO versions of the application, the get-mail and display-message functions are performed by the same program (CSQ4TVDP2). In the CICS version of the application, these functions are performed by separate programs (CSQ4CVD2 and CSQ4CVD3).

The Mail Awaiting panel (CSQ4VDP2 for TSO, VD2 for CICS ; see Figure 97 for an example) shows all the messages that are on the user's mail queue. To create this list, the program uses the MQGET call to browse all the messages on the queue, saving information about each one. In addition to the information displayed, the program records the MsgId and CorrelId of each message.

```
----- IBM MQ for z/OS Sample Programs ----- ROW 16 OF 29
COMMAND ==> Scroll ==> PAGE
USERID - NTSFV02
Mail Manager System QMGR - VC4
Mail Awaiting

Msg Mail Date Time
No From Sent Sent
16
16 Deleted
17 JOHNNJ 01/06/1993 12:52:02
18 JOHNNJ 01/06/1993 12:52:02
19 JOHNNJ 01/06/1993 12:52:03
20 JOHNNJ 01/06/1993 12:52:03
21 JOHNNJ 01/06/1993 12:52:03
22 JOHNNJ 01/06/1993 12:52:04
23 JOHNNJ 01/06/1993 12:52:04
24 JOHNNJ 01/06/1993 12:52:04
25 JOHNNJ 01/06/1993 12:52:05
26 JOHNNJ 01/06/1993 12:52:05
27 JOHNNJ 01/06/1993 12:52:05
28 JOHNNJ 01/06/1993 12:52:06
29 JOHNNJ 01/06/1993 12:52:06
```

Figure 97. Example of a panel showing a list of waiting messages

From the Mail Awaiting panel the user can select one message and display the contents of the message (see Figure 98 on page 603 for an example). The program uses the MQGET call to remove this message from the queue, using the MsgId and CorrelId that the program noted when it browsed all the messages. This MQGET call is performed using the MQGMO\_SYNCPOINT option. The program displays the contents of the message, then declares a syncpoint: this commits the MQGET call, so the message now no longer exists.



## Nickname program on z/OS:

When the user defines a nickname, the nickname program (CSQ4TVD5 for TSO, CSQ4CVD5 for CICS ) creates a queue that has the nickname as part of its name.

The program does this by putting a message on the system-command input queue. The message contains the IBM MQ for z/OS command DEFINE QALIAS or DEFINE QREMOTE. The type of queue that the program creates depends on how the user has filled the fields of the Create Nickname panel (CSQ4VDP5 for TSO, VD5 for CICS ):

- If the user has specified only a user name, or the queue manager name is the same as that of the queue manager to which the Mail Manager is connected, the program creates an alias queue.
- If the user has specified both a user name and a queue manager name (and the queue manager is not the one to which the Mail Manager is connected), the program creates a local definition of a remote queue. The program does not check the existence of the queue to which this definition resolves, or even that the remote queue manager exists.

The program also creates a temporary dynamic queue to handle replies from the system-command input queue.

If the queue manager cannot create the nickname queue for a reason that the program expects (for example, the queue already exists), the program displays its own error message. If the queue manager cannot create the queue for a reason that the program does not expect, the program displays up to two of the error messages that are returned to the program by the command server.

**Note:** For each nickname, the nickname program creates only an alias queue or a local definition of a remote queue. The local queues to which these queue names resolve are created only when the user ID that is contained in the nickname is used to start the Mail Manager application.

## The Credit Check sample on z/OS:

The Credit Check sample application is a suite of programs that demonstrates how to use many of the features provided by IBM MQ for z/OS. It shows how the many component programs of an application can pass messages to each other using message queuing techniques.

The sample can run as a stand-alone CICS application. However, to demonstrate how to design a message queuing application that uses the facilities provided by both the CICS and IMS environments, one module is also supplied as an IMS batch message processing program. This extension to the sample is described in “The IMS extension to the Credit Check sample on z/OS” on page 616.

You can also run the sample on more than one queue manager, and send messages between each instance of the application. To do so, see “The Credit Check sample with multiple queue managers on z/OS” on page 616.

The CICS programs are delivered in C and COBOL. The single IMS program is delivered only in C. The supplied data sets are shown in Table 95 on page 583 and Table 97 on page 584.

The application demonstrates a method of assessing the risk when bank customers ask for loans. The application shows how a bank could work in two ways to process loan requests:

- When dealing directly with a customer, bank staff want immediate access to account and credit-risk information.
- When dealing with written applications, bank staff can submit a series of requests for account and credit-risk information, and deal with the replies at a later time.

The financial and security details in the application have been kept simple so that the message queuing techniques are clear.

To prepare and run the Credit Check sample, perform the following steps:

1. Create the VSAM data set that holds information about some example accounts. Do this by editing and running the JCL supplied in data set CSQ4FILE.
2. Perform the steps in “Preparing the sample applications for the CICS environment on z/OS” on page 581. (The additional steps that you must perform if you want to use the IMS extension to the sample are described in “The IMS extension to the Credit Check sample on z/OS” on page 616.)
3. Start the CKTI trigger monitor (supplied with IBM MQ for z/OS ) against queue CSQ4SAMP.INITIATION.QUEUE, using the CICS transaction CKQC.
4. To start the application, start your CICS system and use the transaction MVB1.
5. Select **Immediate** or **Batch** inquiry from the first panel.

The immediate and batch inquiry panels are similar; Figure 99 shows the Immediate Inquiry panel.

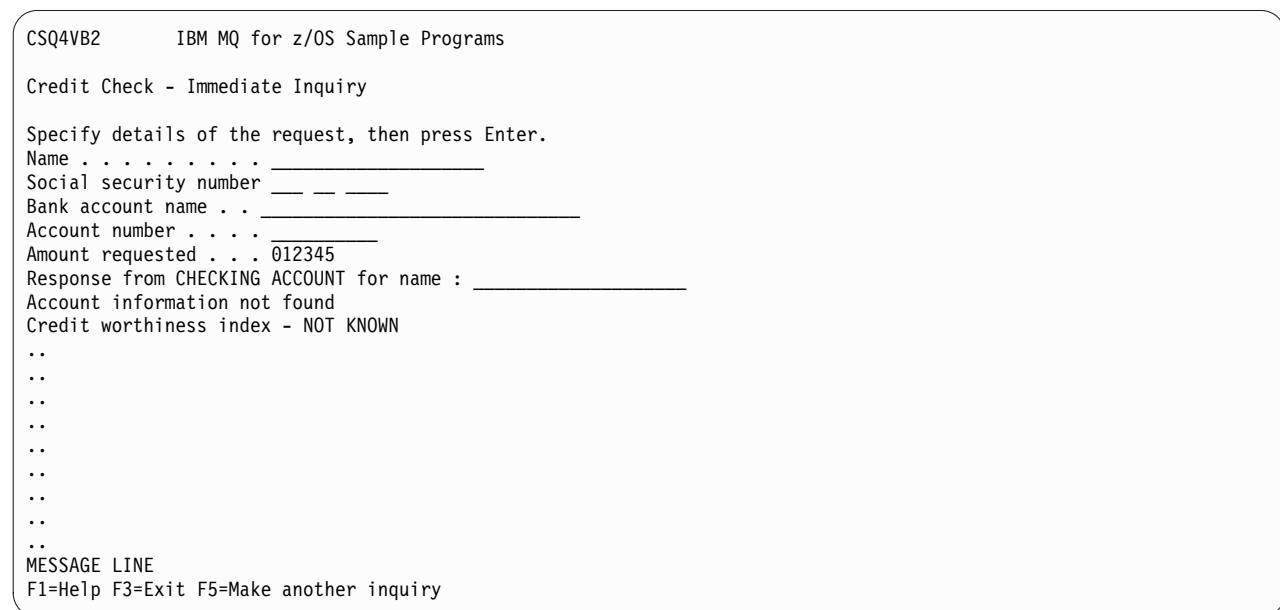


Figure 99. Immediate Inquiry panel for the Credit Check sample application

6. Enter an account number and loan amount in the appropriate fields. See “Entering information in the inquiry panels” for guidance on what information to enter in these fields.

### Entering information in the inquiry panels


The Credit Check sample application checks that the data you enter in the **Amount requested** field of the inquiry panels is in the form of integers.

If you enter one of the following account numbers, the application finds the appropriate account name, average account balance, and credit worthiness index in the VSAM data set CSQ4BAQ:

- 222222222
- 3111234329
- 3256478962
- 333333333
- 3501676212
- 3696879656
- 444444444

- 5555555555
- 6666666666
- 7777777777

You can enter any, or no, information in the other fields. The application retains any information that you enter and returns the same information in the reports that it generates.

*Design of the Credit Check sample on z/OS:* 

This section describes the design of each of the programs that make up the Credit Check sample application.

For more information about some of the techniques that were considered during the design of the application, see “Design considerations for the Credit check sample on z/OS” on page 614.

Figure 100 on page 607 shows the programs that make up the application, and also the queues that these programs serve. In this figure, the prefix CSQ4SAMP has been omitted from all the queue names to make the figure easier to understand.

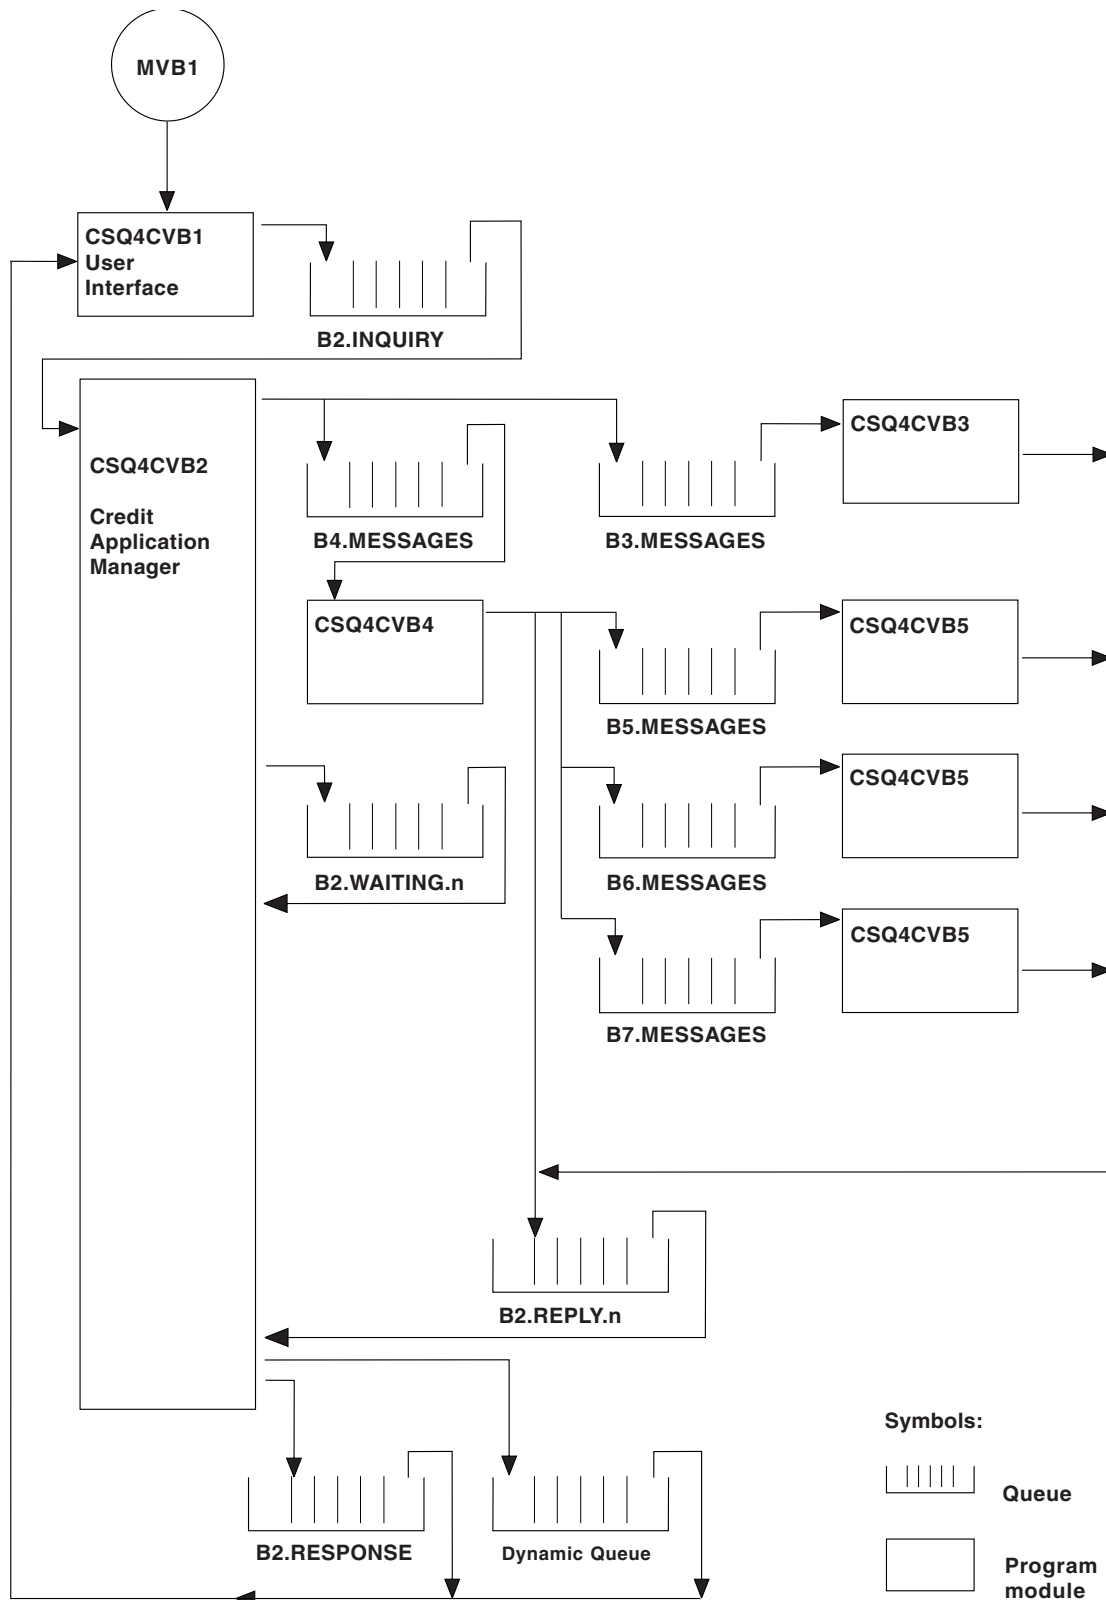


Figure 100. Programs and queues for the Credit Check sample application (COBOL programs only). In the sample application, the queue names shown in this figure have the prefix CSQ4SAMP.

User interface program (CSQ4CVB1) on z/OS: 

When you start the conversational-mode CICS transaction MVB1, this starts the user interface program for the application.


This program puts inquiry messages on queue CSQ4SAMP.B2.INQUIRY and gets replies to those inquiries from a reply-to queue that it specifies when it makes the inquiry. From the user interface you can submit either immediate or batch inquiries:

- For immediate inquiries, the program creates a temporary dynamic queue that it uses as a reply-to queue. This means that each inquiry has its own reply-to queue.
- For batch inquiries, the user-interface program gets replies from the queue CSQ4SAMP.B2.RESPONSE. For simplicity, the program gets replies for all its inquiries from this one reply-to queue. It is easy to see that a bank might want to use a separate reply-to queue for each user of MVB1, so that they could each see replies to only those inquiries that they had initiated.

Important differences between the properties of messages used in the application when in batch and immediate mode are:

- For batch working, the messages have a low priority, so they are processed after any loan requests that are entered in immediate mode. Also, the messages are persistent, so they are recovered if the application or the queue manager has to restart.
- For immediate working, the messages have a high priority, so they are processed before any loan requests that are entered in batch mode. Also, messages are not persistent so they are discarded if the application or the queue manager has to restart.

However, in all cases, the properties of loan request messages are propagated throughout the application. So, for example, all messages that result from a high-priority request will also have a high priority.

Credit application manager (CSQ4CVB2) on z/OS: 

The Credit Application Manager (CAM) program performs most of the processing for the Credit Check application.

The CAM is started by the CKTI trigger monitor (supplied with IBM MQ for z/OS) when a trigger event occurs on either queue CSQ4SAMP.B2.INQUIRY or queue CSQ4SAMP.B2.REPLY.*n*, where *n* is an integer that identifies one of a set of reply queues. The trigger message contains data that includes the name of the queue on which the trigger event occurred.

The CAM uses queues with names of the form CSQ4SAMP.B2.WAITING.*n* to store information about inquiries that it is processing. The queues are named so that they are each paired with a reply-to queue; for example, queue CSQ4SAMP.B2.WAITING.3 contains the input data for a particular inquiry, and queue CSQ4SAMP.B2.REPLY.3 contains a set of reply messages (from programs that query databases) all relating to that same inquiry. To understand the reasons behind this design, see “Separate inquiry and reply queues in the CAM” on page 614.

## Startup logic

If the trigger event occurs on queue CSQ4SAMP.B2.INQUIRY, the CAM opens the queue for shared access. It then tries to open each reply queue until a free one is found. If it cannot find a free reply queue, the CAM logs the fact and terminates normally.

If the trigger event occurs on queue CSQ4SAMP.B2.REPLY.*n*, the CAM opens the queue for exclusive access. If the return code reports that the object is already in use, the CAM terminates normally. If any other error occurs, the CAM logs the error and terminates. The CAM opens the corresponding waiting queue and the inquiry queue, then starts getting and processing messages. From the waiting queue, the CAM recovers details of partially-completed inquiries.



For the sake of simplicity in this sample, the names of the queues used are held in the program. In a business environment, the queue names would probably be held in a file accessed by the program.

### Getting a message from the enquiry queue

The CAM first attempts to get a message from the inquiry queue using the MQGET call with the MQGMO\_SET\_SIGNAL option. If a message is available immediately, the message is processed; if no message is available, a signal is set.

The CAM then attempts to get a message from the reply queue, again using the MQGET call with the same option. If a message is available immediately, the message is processed; otherwise a signal is set.

When both signals are set, the program waits until one of the signals is posted. If a signal is posted to indicate that a message is available, the message is retrieved and processed. If the signal expires or the queue manager is terminating, the program terminates.

### Processing the message retrieved by the CAM

A message retrieved by the CAM can be one of four types:

- An inquiry message
- A reply message
- A propagation message
- An unexpected or unwanted message

The CAM processes these messages as described in “Processing the message retrieved by the CAM on z/OS” on page 610.

### Sending an answer


When the CAM has received all the replies it is expecting for an inquiry, it processes the replies and creates a single response message. It consolidates into one message all the data from all reply messages that have the same `CorrelId`. This response is put on the reply-to queue specified in the original loan request. The response message is put within the same unit of work that contains the retrieval of the final reply message. This is to simplify recovery by ensuring that there is never a completed message on queue `CSQ4SAMP.B2.WAITING.n`.

### Recovery of partially-completed inquiries

The CAM copies onto queue `CSQ4SAMP.B2.WAITING.n` all the messages that it receives. It sets the fields of the message descriptor like this:


- *Priority* is determined by the type of message:
  - For request messages, priority = 3
  - For datagrams, priority = 2
  - For reply messages, priority = 1
- *CorrelId* is set to the *MsgId* of the loan request message
- Other MQMD fields are copied from those of the received message

When an inquiry has been completed, the messages for a specific inquiry are removed from the waiting queue during answer processing. Therefore, at any time, the waiting queue contains all messages relevant to in-progress inquiries. These messages are used to recover details of in-progress inquiries if the program has to restart. The different priorities are set so that inquiry messages are recovered before propagations or reply messages.

Startup logic: 


If the trigger event occurs on queue CSQ4SAMP.B2.INQUIRY, the CAM opens the queue for shared access. It then tries to open each reply queue until a free one is found. If it cannot find a free reply queue, the CAM logs the fact and terminates normally.

This information has moved. See “Credit application manager (CSQ4CVB2) on z/OS” on page 608.

Getting a message from the enquiry queue: 

The CAM first attempts to get a message from the inquiry queue using the MQGET call with the MQGMO\_SET\_SIGNAL option. If a message is available immediately, the message is processed; if no message is available, a signal is set.

This information has moved. See “Credit application manager (CSQ4CVB2) on z/OS” on page 608.

Processing the message retrieved by the CAM on z/OS: 

A message retrieved by the Credit Application Manager (CAM) can be one of four types. The way in which the CAM processes a message depends on its type.

A message retrieved by the CAM can be one of four types:

- An inquiry message
- A reply message
- A propagation message
- An unexpected or unwanted message

The CAM processes these messages as follows:

#### **Inquiry message**

Inquiry messages come from the user interface program. It creates an inquiry message for each loan request.

For all loan requests, the CAM requests the average balance of the customer's checking account. It does this by putting a request message on alias queue CSQ4SAMP.B2.OUTPUT.ALIAS. This queue name resolves to queue CSQ4SAMP.B3.MESSAGES, which is processed by the checking-account program, CSQ4CVB3. When the CAM puts a message on this alias queue, it specifies the appropriate CSQ4SAMP.B2.REPLY.n queue for the reply-to queue. An alias queue is used here so that program CSQ4CVB3 can easily be replaced by another program that processes a base queue of a different name. To do this, you redefine the alias queue so that its name resolves to the new queue. Also, you could assign differing access authorities to the alias queue and to the base queue.

If a user requests a loan that is larger than 10000 units, the CAM initiates checks on other databases as well. It does this by putting a request message on queue CSQ4SAMP.B4.MESSAGES, which is processed by the distribution program, CSQ4CVB4. The process serving this queue propagates the message to queues served by programs that have access to other records such as credit card history, savings accounts, and mortgage payments. The data from these programs is returned to the reply-to queue specified in the put operation. Additionally, a propagation message is sent to the reply-to queue by this program to specify how many propagation messages have been sent.

In a business environment, the distribution program would probably reformat the data provided to match the format required by each of the other types of bank account.

Any of the queues referred to can be on a remote system.

For each inquiry message, the CAM initiates an entry in the memory-resident Inquiry Record Table (IRT). This record contains:

- The MsgId of the inquiry message
- In the ReplyExp field, the number of responses expected (equal to the number of messages sent)
- In the ReplyRec field, the number of replies received (zero at this stage)
- In the PropsOut field, an indication of whether a propagation message is expected

The CAM copies the inquiry message onto the waiting queue with:

- Priority set to 3
- CorrelId set to the MsgId of the inquiry message
- The other message-descriptor fields set to those of the inquiry message

### **Propagation message**

A propagation message contains the number of queues to which the distribution program has forwarded the inquiry. The message is processed as follows:

1. Add to the ReplyExp field of the appropriate record in the IRT the number of messages sent. This information is in the message.
2. Increment by 1 the ReplyRec field of the record in the IRT.
3. Decrement by 1 the PropsOut field of the record in the IRT.
4. Copy the message onto the waiting queue. The CAM sets the Priority to 2 and the other fields of the message descriptor to those of the propagation message.

### **Reply message**


A reply message contains the response to one of the requests to the checking-account program or to one of the agency-query programs. Reply messages are processed as follows:

1. Increment by 1 the ReplyRec field of the record in the IRT.
2. Copy the message onto the waiting queue with Priority set to 1 and the other fields of the message descriptor set to those of the reply message.
3. If ReplyRec = ReplyExp, and PropsOut = 0, set the MsgComplete flag.

### **Other messages**


The application does not expect other messages. However, the application might receive messages broadcast by the system, or reply messages with a unknown CorrelIds.

The CAM puts these messages on queue CSQ4SAMP.DEAD.QUEUE, where they can be examined. If this put operation fails, the message is lost and the program continues. For more information about the design of this part of the program, see "How the sample handles unexpected messages" on page 614.

*Sending an answer:*  z/OS

When the CAM has received all the replies it is expecting for an inquiry, it processes the replies and creates a single response message.

This information has moved. See “Credit application manager (CSQ4CVB2) on z/OS” on page 608.

*Recovery of partially-completed inquiries:*  z/OS

The CAM copies onto queue CSQ4SAMP.B2.WAITING.n all the messages that it receives and sets the fields of the message descriptor.

This information has moved. See “Credit application manager (CSQ4CVB2) on z/OS” on page 608.

*Checking-account program (CSQ4CVB3) on z/OS:*  z/OS


The checking-account program is started by a trigger event on queue CSQ4SAMP.B3.MESSAGES. After it has opened the queue, this program gets a message from the queue using the MQGET call with the wait option, and with the wait interval set to 30 seconds.

The program searches VSAM data set CSQ4BAQ for the account number in the loan request message. It retrieves the corresponding account name, average balance, and credit worthiness index, or notes that the account number is not in the data set.

The program then puts a reply message (using the MQPUT1 call) on the reply-to queue named in the loan request message. For this reply message, the program:

- Copies the CorrelId of the loan request message
- Uses the MQPMO\_PASS\_IDENTITY\_CONTEXT option

The program continues to get messages from the queue until the wait interval expires.

*Distribution program (CSQ4CVB4) on z/OS:*  z/OS

The distribution program is started by a trigger event on queue CSQ4SAMP.B4.MESSAGES.

To simulate the distribution of the loan request to other agencies that have access to records such as credit card history, savings accounts, and mortgage payments, the program puts a copy of the same message on all the queues in the namelist CSQ4SAMP.B4.NAMELIST. There are three of these queues, with names of the form CSQ4SAMP.B n.MESSAGES, where n is 5, 6, or 7. In a business application, the agencies could be at separate locations, so these queues could be remote queues. If you want to modify the sample application to show this, see “The Credit Check sample with multiple queue managers on z/OS” on page 616.

The distribution program performs the following steps:

1. From the namelist, gets the names of the queues that the program is to use. The program does this by using the MQINQ call to inquire about the attributes of the namelist object.
2. Opens these queues and also CSQ4SAMP.B4.MESSAGES.
3. Performs the following loop until there are no more messages on queue CSQ4SAMP.B4.MESSAGES:
  - a. Get a message using the MQGET call with the wait option, and with the wait interval set to 30 seconds.

- b. Put a message on each queue listed in the namelist, specifying the name of the appropriate CSQ4SAMP.B2.REPLY.n queue for the reply-to queue. The program copies the *CorrelId* of the loan request message to these copy messages, and it uses the MQPMO\_PASS\_IDENTITY\_CONTEXT option on the MQPUT call.
- c. Send a datagram message to queue CSQ4SAMP.B2.REPLY.n to show how many messages it has successfully put.
- d. Declare a syncpoint.

Agency-query program (CSQ4CVB5/CSQ4CCB5) on z/OS: 

The agency-query program is supplied as both a COBOL program and a C program. Both programs have the same design. This shows that programs of different types can easily coexist within an IBM MQ application, and that the program modules that make up such an application can easily be replaced.

An instance of the program is started by a trigger event on any of these queues:

- For the COBOL program (CSQ4CVB5):
  - CSQ4SAMP.B5.MESSAGES
  - CSQ4SAMP.B6.MESSAGES
  - CSQ4SAMP.B7.MESSAGES
- For the C program (CSQ4CCB5), queue CSQ4SAMP.B8.MESSAGES

**Note:** If you want to use the C program, you must alter the definition of the namelist CSQ4SAMP.B4.NAMELIST to replace the queue CSQ4SAMP.B7.MESSAGES with CSQ4SAMP.B8.MESSAGES. To do this, you can use any one of:

- The IBM MQ for z/OS operations and control panels
- The ALTER NAMELIST command
- The CSQUTIL utility

After it has opened the appropriate queue, this program gets a message from the queue using the MQGET call with the wait option, and with the wait interval set to 30 seconds.

The program simulates the search of an agency's database by searching the VSAM data set CSQ4BAQ for the account number that was passed in the loan request message. It then builds a reply that includes the name of the queue that it is serving and a creditworthiness index. To simplify the processing, the creditworthiness index is selected at random.

When putting the reply message, the program uses the MQPUT1 call and:

- Copies the CorrelId of the loan request message
- Uses the MQPMO\_PASS\_IDENTITY\_CONTEXT option

The program sends the reply message to the reply-to queue named in the loan request message. (The name of the queue manager that owns the reply-to queue is also specified in the loan request message.)

Design considerations for the Credit Check sample.

This topic contains information about:

- “Separate inquiry and reply queues in the CAM”
- “How the sample handles errors”
- “How the sample handles unexpected messages”
- “How the sample uses syncpoints” on page 615
- “How the sample uses message context information” on page 615
- “Use of message and correlation identifiers in the CAM” on page 616

### **Separate inquiry and reply queues in the CAM**

The application could use a single queue for both inquiries and replies, but it was designed to use separate queues for the following reasons:

- When the program is handling the maximum number of inquiries, further inquiries can be left on the queue. If a single queue is being used, this would have to be taken off the queue and stored elsewhere.
- Other instances of the CAM could be started automatically to service the same inquiry queue if message traffic was high enough to warrant it. But the program must track in-progress inquiries, and to do this, it must get back all replies to inquiries it has initiated. If only one queue is used, the program would have to browse the messages to see if they were for this program or for another. This would make the operation much less efficient.

The application can support multiple CAMs and can recover in-progress inquiries effectively by using paired reply-to and waiting queues.

- The program can wait on multiple queues effectively by using signaling.

### **How the sample handles errors**

The user interface program handles errors by reporting them directly to the user.

The other programs do not have user interfaces, so they have to handle errors in other ways. Also, in many situations (for example, if an MQGET call fails) these other programs do not know the identity of the user of the application.

The other programs put error messages on a CICS temporary storage queue called CSQ4SAMP. You can browse this queue using the CICS-supplied transaction CEBR. The programs also write error messages to the CICS CSML log.

### **How the sample handles unexpected messages**

When you design a message-queuing application, you must decide how to handle messages that arrive on a queue unexpectedly.

The two basic choices are:

- The application does no more work until it has processed the unexpected message. This probably means that the application notifies an operator, terminates itself, and ensures that it is not restarted automatically (it can do this by setting triggering off). This choice means that all processing for the application can be halted by a single unexpected message, and the intervention of an operator is required to restart the application.
- The application removes the message from the queue it is serving, puts the message in another location, and continues processing. The best place to put this message is on the system dead-letter queue.

If you choose the second option:

- An operator, or another program, should examine the messages that are put on the dead-letter queue to find out where the messages are coming from.
- An unexpected message is lost if it cannot be put on the dead-letter queue.
- A long unexpected message is truncated if it is longer than the limit for messages on the dead-letter queue, or longer than the buffer size in the program.

To ensure that the application smoothly handles all inquiries with minimal effect from outside activities, the Credit Check sample application uses the second option. To allow you to keep the sample separate from other applications that use the same queue manager, the Credit Check sample does not use the system dead-letter queue; instead, it uses its own dead-letter queue. This queue is named CSQ4SAMP.DEAD.QUEUE. The sample truncates any messages that are longer than the buffer area provided for the sample programs. You can use the Browse sample application to browse messages on this queue, or use the Print Message sample application to print the messages together with their message descriptors.

However, if you extend the sample to run across more than one queue manager, unexpected messages, or messages that cannot be delivered, could be put on the system dead-letter queue by the queue manager.

### **How the sample uses syncpoints**

The programs in the Credit Check sample application declare syncpoints to ensure that:

- Only one reply message is sent in response to each expected message
- Multiple copies of unexpected messages are never put on the sample's dead-letter queue
- The CAM can recover the state of all partially completed inquiries by getting persistent messages from its waiting queue

To achieve this, a single unit of work is used to cover the getting of a message, the processing of that message, and any subsequent put operations.

### **How the sample uses message context information**

When the user interface program (CSQ4CVB1) sends messages, it uses the MQPMO\_DEFAULT\_CONTEXT option. This means that the queue manager generates both identity and origin context information. The queue manager gets this information from the transaction that started the program (MVB1) and from the user ID that started the transaction.

When the CAM sends inquiry messages, it uses the MQPMO\_PASS\_IDENTITY\_CONTEXT option. This means that the identity context information of the message being put is copied from the identity context of the original inquiry message. With this option, origin context information is generated by the queue manager.

When the CAM sends reply messages, it uses the MQPMO\_ALTERNATE\_USER\_AUTHORITY option. This causes the queue manager to use an alternate user ID for its security check when the CAM opens a reply-to queue. The CAM uses the user ID of the submitter of the original inquiry message. This means that users are allowed to see replies to only those inquiries that they have originated. The alternate user ID is obtained from the identity context information in the message descriptor of the original inquiry message.


When the query programs (CSQ4CVB3/4/5) send reply messages, they use the MQPMO\_PASS\_IDENTITY\_CONTEXT option. This means that the identity context information of the message being put is copied from the identity context of the original inquiry message. With this option, origin context information is generated by the queue manager.

**Note:** The user ID associated with the MVB3/4/5 transactions requires access to the B2.REPLY.n queues. These user IDs might not be the same as those associated with the request being processed. To get around this possible security exposure, the query programs could use the MQPMO\_ALTERNATE\_USER\_AUTHORITY option when putting their replies. This would mean that each individual user of MVB1 needs authority to open the B2.REPLY.n queues.

### Use of message and correlation identifiers in the CAM

The application has to monitor the progress of all the live inquiries it is processing at any one time. To do this it uses the unique message identifier of each loan request message to associate all the information that it has about each inquiry.

The CAM copies the MsgId of the inquiry message into the CorrelId of all the request messages it sends for that inquiry. The other programs in the sample (CSQ4CVB3 - 5) copy the CorrelId of each message that they receive into the CorrelId of their reply message.

*The Credit Check sample with multiple queue managers on z/OS:* 

You can use the Credit Check sample application to demonstrate distributed queuing by installing the sample on two queue managers and CICS systems (with each queue manager connected to a different CICS system).

When the sample program is installed, and the trigger monitor (CKTI) is running on each system, you need to:

1. Set up the communication link between the two queue managers. For information on how to do this, see *Configuring distributed queuing*.
2. On one queue manager, create a local definition for each of the remote queues (on the other queue manager) that you want to use. These queues can be any of CSQ4SAMP.B *n*.MESSAGES, where *n* is 3, 5, 6, or 7. (These are the queues that are served by the checking-account program and the agency-query program.) For information on how to do this, see *DEFINE QREMOTE* and *DEFINE* queues.
3. Change the definition of the namelist (CSQ4SAMP.B4.NAMELIST) so that it contains the names of the remote queues that you want to use. For information on how to do this, see *DEFINE NAMELIST*.

*The IMS extension to the Credit Check sample on z/OS:* 

A version of the checking-account program is supplied as an IMS batch message processing (BMP) program. It is written in the C language.

The program performs the same function as the CICS version, except that to obtain the account information, the program reads an IMS database instead of a VSAM file. If you replace the CICS version of the checking-account program with the IMS version, you see no difference in the method of using the application.

To prepare and run the IMS version you must:

1. Follow the steps in “Preparing and running the Credit Check sample on z/OS” on page 605.
2. Follow the steps in “Preparing the sample application for the IMS environment on z/OS” on page 584.
3. Alter the definition of the alias queue CSQ4SAMP.B2.OUTPUT.ALIAS to resolve to queue CSQ4SAMP.B3.IMS.MESSAGES (instead of CSQ4SAMP.B3.MESSAGES). To do this, you can use one of:
  - The IBM MQ for z/OS operations and control panels
  - The ALTER QALIAS command .



Another way of using the IMS checking-account program is to make it serve one of the queues that receives messages from the distribution program. In the delivered form of the Credit Check sample application, there are three of these queues (B5/6/7.MESSAGES), all served by the agency-query program. This program searches a VSAM data set. To compare the use of the VSAM data set and the IMS database, you could make the IMS checking-account program serve one of these queues instead. To do this, you must alter the definition of the namelist CSQ4SAMP.B4.NAMELIST to replace one of the CSQ4SAMP.B *n*.MESSAGES queues with the CSQ4SAMP.B3.IMS.MESSAGES queue. You can use one of:

- The IBM MQ for z/OS operations and control panels
- The ALTER NAMELIST command.

You can then run the sample from CICS transaction MVB1. The user sees no difference in operation or response. The IMS BMP stops either after receiving a stop message or after being inactive for five minutes.

### Design of the IMS checking-account program (CSQ4ICB3)

This program runs as a BMP. Start the program using its JCL before any IBM MQ messages are sent to it.

The program searches an IMS database for the account number in the loan request messages. It retrieves the corresponding account name, average balance, and credit worthiness index.

The program sends the results of the database search to the reply-to queue named in the IBM MQ message being processed. The message returned appends the account type and the results of the search to the message received so that the transaction building the response can confirm that the correct query is being processed. The message is in the form of three 79-character groups, as follows:

```
'Response from CHECKING ACCOUNT for name : JONES J B'
' Opened 870530, 3-month average balance = 000012.57'
' Credit worthiness index - BBB'
```

When running as a message-oriented BMP, the program drains the IMS message queue, then reads messages from the IBM MQ for z/OS queue and processes them. No information is received from the IMS message queue. The program reconnects to the queue manager after each checkpoint because the handles have been closed.

When running in a batch-oriented BMP, the program continues to be connected to the queue manager after each checkpoint because the handles are not closed.

### The Message Handler sample on z/OS:

The Message Handler sample TSO application allows you to browse, forward, and delete messages on a queue. The sample is available in C and COBOL.

### Preparing and running the sample


Follow these steps:

1. Prepare the sample as described in “Preparing sample applications for the TSO environment on z/OS” on page 579.
2. Tailor the CLIST (CSQ4RCH1) provided for the sample to define the location of the panels, the location of the message file, and the location of the load modules.

You can use CLIST CSQ4RCH1 to run both the C and the COBOL version of the sample. The supplied version of CSQ4RCH1 runs the C version, and contains instructions on the tailoring necessary for the COBOL version.

### Note:

1. There are no sample queue definitions provided with the sample.
2. VS COBOL II does not support multitasking with ISPF, so do not use the Message Handler sample application on both sides of a split screen. If you do, the results are unpredictable.

Using the Message Handler sample on z/OS: 

Having installed the sample and invoked it from the tailored CLIST CSQ4RCH1, the screen shown in Figure 101 is displayed.

```

----- IBM MQ for z/OS -- Samples -----
COMMAND ==>
User Id : JOHNJ

Enter information. Press ENTER :

Queue Manager Name : _____ :
Queue Name : _____ :

F1=HELP F2=SPLIT F3=END F4=RETURN F5=RFIND F6=RCHANGE
F7=UP F8=DOWN F9=SWAP F10=LEFT F11=RIGHT F12=RETRIEVE

```

Figure 101. Initial screen for Message Handler sample

Enter the queue manager and queue name to be viewed (case sensitive) and the message list screen is displayed (see Figure 102 ).

```

----- IBM MQ for z/OS -- Samples ----- Row 1 to 4 of 4
COMMAND ==>
Queue Manager : VM03 :
Queue : MQEI.IMS.BRIDGE.QUEUE :
Message number 01 of 04

Msg Put Date Put Time Format User Put Application
No MM/DD/YYYY HH:MM:SS Name Identifier Type Name
01 10/16/1998 13:51:19 MQIMS NTSFV02 00000002 NTSFV02A
02 10/16/1998 13:55:45 MQIMS JOHNJ 00000011 EDIT\CLASSES\BIN\PROGTS
03 10/16/1998 13:54:01 MQIMS NTSFV02 00000002 NTSFV02B
04 10/16/1998 13:57:22 MQIMS johnj 00000011 EDIT\CLASSES\BIN\PROGTS
***** Bottom of data *****

```

Figure 102. Message list screen for Message Handler sample

This screen shows the first 99 messages on the queue and, for each, shows the following fields:

**Msg No**

Message number

**Put Date MM/DD/YYYY**

Date that the message was put on the queue (GMT)

**Put Time HH:MM:SS**

Time that the message was put on the queue (GMT)

**Format Name**

MQMD.Format field

**User Identifier**

MQMD.UserIdentifier field

**Put Application Type**

MQMD.PutApplType field

**Put Application Name**

MQMD.PutApplName field

The total number of messages on the queue is also displayed.

From this screen a message can be chosen, by number not by cursor position, and then displayed. For an example, see Figure 103 on page 620.

```

----- IBM MQ for z/OS -- Samples ----- Row 1 to 35 of 35
COMMAND ==>

Queue Manager : VM03 :
Queue : MQEI.IMS.BRIDGE.QUEUE :
Forward to Q Mgr : VM03 :
Forward to Queue : QL.TEST.ISCRES1 :

Action : _ : (D)elete (F)orward

Message Content :

Message Descriptor
StrucId : `MD `
Version : 000000001
Report : 000000000
MsgType : 000000001
Expiry : -00000001
Feedback : 000000000
Encoding : 000000785
CodedCharSetId : 000000500
Format : `MQIMS `
Priority : 000000000
Persistence : 000000001
MsgId : `C3E2D840E5D4F0F34040404040404040AF6B30F0A89B7605`X
CorrelId : `00`X
BackoutCount : 000000000
ReplyToQ : `QL.TEST.ISCRES1
ReplyToQMgr : `VM03
UserIdentifier : `NTSFV02
AccountingToken :
`06F2F5F5F3F0F100`X
ApplIdentityData : `
PutApplType : 000000002
PutApplName : `NTSFV02A
PutDate : `19971016`
PutTime : `13511903`
ApplOriginData : `

Message Buffer : 108 byte(s)
00000000 : C9C9 C840 0000 0001 0000 0054 0000 0311 `IIH`
00000010 : 0000 0000 4040 4040 4040 4040 0000 0000 `.... `
00000020 : 4040 4040 4040 4040 4040 4040 4040 4040 `.....`
00000030 : 4040 4040 4040 4040 4040 4040 4040 4040 `.....`
00000040 : 0000 0000 0000 0000 0000 0000 0000 0000 `.....`
00000050 : 40F1 C300 0018 0000 C9C1 D7D4 C4C9 F2F8 `1C.....IAPMDI28`
00000060 : 40C8 C5D3 D3D6 40E6 D6D9 D3C4 `HELLO WORLD `
***** Bottom of data *****

```

Figure 103. Chosen message is displayed

Once the message has been displayed it can be deleted, left on the queue, or forwarded to another queue. The Forward to Q Mgr and Forward to Queue fields are initialized with values from the MQMD, these can be changed before forwarding the message.

The sample design allows only messages with unique MsgId / CorrelId combinations to be selected and displayed, because the message is retrieved using the MsgId and CorrelId as the key. If the key is not unique the sample cannot retrieve the chosen message with certainty.

**Note:** When you use the SCSQCLST(CSQ4RCH1) sample to browse messages, each invocation causes the backout count of the message to increase. If you want to change the behavior of this sample, copy the sample and modify the contents as necessary. You should be aware that other applications that rely on this backout count can be influenced by this increasing count.

This topic describes the design of each of the programs that make up the Message Handler sample application.

### **Object validation program**

This requests a valid queue and queue manager name.

If you do not specify a queue manager name, the default queue manager is used, if available. Only local queues can be used; an MQINQ is issued to check that the queue type and an error is reported if the queue is not local. If the queue is not opened successfully, or the MQGET call is inhibited on the queue, error messages are returned indicating the CompCode and Reason return code.

### **Message list program**

This displays a list of messages on a queue with information about them such as the putdate, puttime, and the message format.

The maximum number of messages stored in the list is 99. If there are more messages on the queue than this, the current queue depth is also displayed. To choose a message for display, type the message number into the entry field (the default is 01). If your entry is not valid, you receive an appropriate error message.

### **Message content program**

This displays message content.

The content is formatted and split into two parts:

1. Message descriptor
2. Message buffer

The message descriptor shows the contents of each field on a separate line.

The message buffer is formatted depending on its contents. If the buffer holds a dead letter header (MQDLH) or a transmission queue header (MQXQH), these are formatted and displayed before the buffer itself.

Before the buffer data is formatted, a title line shows the buffer length of the message in bytes. The maximum buffer size is 32768 bytes, and any message longer than this is truncated. The full size of the buffer is displayed along with a message indicating that only the first 32768 bytes of the message are displayed.

The buffer data is formatted in two ways:

1. After the offset into the buffer is printed, the buffer data is displayed in hexadecimal.
2. The buffer data is then displayed again as EBCDIC values. If any EBCDIC value cannot be printed, it prints a period (.) instead.

You can enter D for delete, or F for forward into the action field. If you choose to forward the message, the forward-to queue and queue manager name must be set correctly. The defaults for these fields are read from the message descriptor ReplyToQ and ReplyToQMgr fields.

If you forward a message, any header block stored in the buffer is stripped. If the message is forwarded successfully, it is removed from the original queue. If you enter invalid actions, error messages are displayed.

An example help panel called CSQ4CHP9 is also available.

### The Asynchronous Put sample on z/OS:

The Asynchronous Put sample program puts messages on a queue using the asynchronous MQPUT call. The sample also retrieves status information using the MQSTAT call.

The Asynchronous Put applications use these MQI calls:

- MQCONN
- MQOPEN
- MQPUT
- MQSTAT
- MQCLOSE
- MQDISC

The sample programs are delivered in the C programming language.

The Asynchronous Put applications run in the batch environment. See Other samples for the batch applications.

This topic also provides information about the design of the Asynchronous Consumption program, and running the CSQ4BCS2 sample.

- “Running the CSQ4BCS2 sample”
- “Design of the Asynchronous Put sample program”

### Running the CSQ4BCS2 sample

This sample program takes up to six parameters:

1. The name of the target queue (required).
2. The name of the queue manager (optional).
3. Open options (optional).
4. Close options (optional).
5. The name of the target queue manager (optional).
6. The name of the dynamic queue (optional).

If a queue manager is not specified, CSQ4BCS2 connects to the default queue manager. Message content is provided through standard input ( **SYSD** ).

There is a sample JCL to run the program, it resides in CSQ4BCSP.

### Design of the Asynchronous Put sample program

The program uses the MQOPEN call with either the output options supplied, or with the MQOO\_OUTPUT and MQOO\_FAIL\_IF QUIESCING options, to open the target queue for putting messages.

If the program cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call. To keep the program simple on this and subsequent MQI calls, default values are used for many of the options.

For each line of input, the program reads the text into a buffer and uses the MQPUT call with MQPMO\_ASYNC\_RESPONSE to create a datagram message containing the text of that line and

asynchronously puts the message on the target queue. The program continues until it reaches the end of the input, or until the MQPUT call fails. If the program reaches the end of the input, it closes the queue using the MQCLOSE call.

The program then issues the MQSTAT call which returns an MQSTS structure, and displays messages containing the number of messages put successfully, the number of messages put with a warning, and the number of failures.

**Note:** To observe what happens when an MQPUT error is detected by the MQSTAT call, set MAXDEPTH on the target queue to a low value.

**The Batch Asynchronous Consumption sample on z/OS:** z/OS

The CSQ4BCS1 sample program is delivered in C, it demonstrates the use of MQCB and MQCTL to consume messages from multiple queues asynchronously.

The Asynchronous Consumption samples run in the batch environment. See Other samples for the batch applications.

There is also a COBOL sample which runs in the CICS environment, see “The CICS Asynchronous Consumption and Publish/Subscribe sample on z/OS” on page 624.

The applications use these MQI calls:


- MQCONN
- MQOPEN
- MQCLOSE
- MQDISC
- MQCB
- MQCTL

This topic also provided information about the following headings:

- “Running the CSQ4BCS1 sample”
- “Design of the Batch Asynchronous Consumption sample program” on page 624

**Running the CSQ4BCS1 sample**

This sample program follows the following syntax:



**Notes:**

- 1 A maximum of 10 queues are supported by this sample. Provide more than one queue name in order to read messages from multiple queues.

There is a sample JCL to run this program, it resides in CSQ4BCSC.

## Design of the Batch Asynchronous Consumption sample program

The sample shows how to read messages from multiple queues in the order of their arrival. This would require more code using synchronous MQGET. With asynchronous consumption, no polling is required, and thread and storage management is performed by IBM MQ. In the sample program, errors are written to the console.

The sample code has the following steps:

1. Define the single message consumption callback function.

```
void MessageConsumer(MQHCONN hConn,
MQMD * pMsgDesc,
MQGMO * pGetMsgOpts,
MQBYTE * Buffer,
MQCBC * pContext)
{ ... }
```

2. Connect to the queue manager.

```
MQCONN(QMName,&Hcon,&CompCode,&CReason);
```

3. Open the input queues, and associate each queue with the MessageConsumer callback function.

```
MQOPEN(Hcon,&od,O_options,&Hobj,&OpenCode,&Reason);
cbd.CallbackFunction = MessageConsumer;
MQCB(Hcon,MQOP_REGISTER,&cbd,Hobj,&md,&gmo,&CompCode,&Reason);
```

cbd.CallbackFunction does not need to be set for each queue; it is an input-only field. You can associate a different callback function with each queue.

4. Start consumption of the messages.

```
MQCTL(Hcon,MQOP_START,&ctlo,&CompCode,&Reason);
```

5. Wait for the user to press Enter, then stop consumption of messages.

```
MQCTL(Hcon,MQOP_STOP,&ctlo,&CompCode,&Reason);
```

6. Finally, disconnect from the queue manager.

```
MQDISC(&Hcon,&CompCode,&Reason);
```

## The CICS Asynchronous Consumption and Publish/Subscribe sample on z/OS:

The Asynchronous Consumption and Publish/Subscribe sample programs demonstrate the use of asynchronous consumption, and publish and subscribe features within CICS.

A *Registration client* program registers three Callback handlers (an event handler, and two message consumers), and starts Asynchronous Consumption. A *Messaging client* program puts messages to a queue, or publishes suitable messages from a CICS console for consumption by the two Message Consumers (CSQ4CVCN and CSQ4CVCT).

To provide runtime control over the behavior of the sample, one of the message consumers can be instructed using the messages it receives, to SUSPEND, RESUME, or DEREGISTER any of the Callback handlers. It can also be used to issue an MQCTL STOP to end Asynchronous Consumption under control. The other message consumer is registered to subscribe to a topic.

Each program issues COBOL DISPLAY statements at appropriate points to display the behavior of the sample.

The applications use these MQI calls:

- MQOPEN
- MQPUT
- MQSUB
- MQGET



- MQCLOSE
- MQCB
- MQCTL

The programs are delivered in the COBOL language. See CICS Asynchronous Consumption and Publish/Subscribe samples for the CICS applications.

This topic also provides the following information:

- “Setup”
- “Registration Client CSQ4CVRG”
- “Event handler CSQ4CVEV”
- “Simple Message Consumer CSQ4CVCN”
- “Control Message Consumer CSQ4CVCT” on page 626
- “Messaging Client CSQ4CVPT” on page 626

## Setup

The names of the Queue and Topic used by the Message Consumers are hardcoded in the Registration and Messaging Client programs.

The Queue, **SAMPLE.CONTROL.QUEUE**, should be defined to the Queue Manager associated with the CICS region before running the sample. The Topic, **News/Media/Movies**, can be defined if required, or it is created at runtime under the default Administrative Object if it does not exist.

CICS programs and transaction definitions can be installed by installing a group: CSQ4SAMP.

## Registration Client CSQ4CVRG

The Registration Client program must be started under the CICS transaction MVRG. It takes no input.

When started, the Registration Client registers the following Callback handlers using MQCB:

- CSQ4CVEV as an Event Handler.
- CSQ4CVCN as a Message Consumer on a topic, **News/Media/Movies**.
- CSQ4CVCT as a Message Consumer on a Queue, **SAMPLE.CONTROL.QUEUE**.

The Registration Client passes a data structure containing the names of all three registered Callback handlers to CSQ4CVCT, together with the object handles associated with the two message consumers.

Having registered the Callback handlers, the Registration Client issues an MQCTL START\_WAIT to start Asynchronous Consumption, and suspend until control is returned to it (for example, by one of the Callback handlers issuing an MQCTL STOP).

## Event handler CSQ4CVEV

When driven, the Event Handler displays a message indicating the call type (for example, START). When driven for IBM MQ reason code CONNECTION\_QUIESCING, the Event Handler issues an MQCTL STOP to end Asynchronous Consumption and return control to the Registration Client.

## Simple Message Consumer CSQ4CVCN

When driven, this Message Consumer displays a message indicating the call type (for example, REGISTER). When driven for the MSG\_REMOVED call type, the Message Consumer retrieves the inbound message and outputs it to the CICS job log.

## Control Message Consumer CSQ4CVCT

When driven, this Message Consumer displays a message indicating the call type (for example, START). When driven for the MSG\_REMOVED call type, the Message Consumer retrieves the inbound message and the data structure passed by the Registration Client. Based on the message content, it issues appropriate MQCB or MQCTL commands to one of the following:

- STOP Asynchronous Consumption (returning control to the Registration Client).
- SUSPEND, RESUME, or Deregister a named Callback handler (including itself).


## Messaging Client CSQ4CVPT

The Messaging Client has two functions:

- It publishes a message to a topic for consumption by the Message Consumer CSQ4CVCN.
- It puts a control message to a queue for consumption by the Control Message Consumer CSQ4CVCT, resulting in a potential change in behavior of the sample.

The Messaging Client program must be started from a CICS console under a CICS transaction, and it takes command line input with the following syntax:

### CSQ4CVPT



### PUBLISH

Publish the Message Text (or a default message) as a Retained Message for consumption by the Simple Message Consumer.

**STOP** Stop Asynchronous Consumption.

### DEREGISTER

Deregister the named Callback handler.

### RESUME

Resume the named Callback handler.

### SUSPEND

Suspend the named Callback handler.

Input fields are positional, and comma-separated. Keywords and Callback Handler names are not case-sensitive.

Examples:

Table 99. Input examples

| Example                        | Description                        |
|--------------------------------|------------------------------------|
| MVMP,PUBLISH,,                 | Publish a default message          |
| MVMP,publish, A short message, | Publish the given text             |
| MVMP,STOP,                     | Stop Asynchronous Consumption      |
| MVMP,DEREGISTER,CSQ4CVEV,      | Deregister the Event Handler       |
| MVMP,resume,csq4cvcn,          | Resume the Simple Message Consumer |

Table 99. Input examples (continued)

| Example                | Description               |
|------------------------|---------------------------|
| MVMP,SUSPEND,CSQ4CVEV, | Suspend the Event Handler |

Where MVMP is the CICS transaction associated with the Messaging Client program CSQ4CVPT.

**Note:**

- Suspending or deregistering all Callback handlers terminates the START\_WAIT issued by the Registration Client, returning control to it, and ending the task.
- Suspending or deregistering the Control Callback Handler has deliberately not been prevented, but it removes the ability to further control the behavior of the sample.

**The Publish/Subscribe sample on z/OS:** 

The Publish/Subscribe sample programs demonstrate the use of the publish and subscribe features in IBM MQ.

There are four C and two COBOL programming language sample programs demonstrating how to program to the IBM MQ Publish/Subscribe interface. The programs are delivered in the C and COBOL language. The applications run in the batch environment; see Publish/Subscribe samples for the batch applications.

There are also COBOL samples that run in the CICS environment; see “The CICS Asynchronous Consumption and Publish/Subscribe sample on z/OS” on page 624.

This topic also provides information about how to run Publish/Subscribe sample programs. These sample programs include:

- “Running the CSQ4BCP1 sample”
- “Running the CSQ4BCP2 sample” on page 628
- “Running the CSQ4BCP3 sample” on page 628
- “Running the CSQ4BCP4 sample” on page 628
- “Running the CSQ4BVP1 sample” on page 629
- “Running the CSQ4BVP2 sample” on page 629

**Running the CSQ4BCP1 sample**

This program is written in C; it publishes messages to a topic. Start one of the subscriber samples before running this program.

This program takes up to four parameters:

1. The name of the target topic string (required).
2. The name of the queue manager (optional).
3. Open options (optional).
4. Close options (optional).

If a queue manager is not specified, CSQ4BCP1 connects to the default queue manager. There is a sample JCL to run the program, it resides in CSQ4BCPP.

Message content is provided through standard input (**SYSDIN DD**).

### Running the CSQ4BCP2 sample

This program is written in C; it subscribes to a topic and prints the messages received.

This program takes up to three parameters:

1. The name of the target topic string (required).
2. The name of the queue manager (optional).
3. MQSD subscription options (optional).

If a queue manager is not specified, CSQ4BCP2 connects to the default queue manager. There is a sample JCL to run the program, it resides in CSQ4BCPS.

### Running the CSQ4BCP3 sample

This program is written in C; it subscribes to a topic using a user-specified destination queue and prints the messages received.

This program takes up to four parameters:

1. The name of the target topic string (required).
2. The name of the destination (required).
3. The name of the queue manager (optional).
4. MQSD subscription options (optional).

If a queue manager is not specified, CSQ4BCP3 connects to the default queue manager. There is a sample JCL to run the program, it resides in CSQ4BCPD.

### Running the CSQ4BCP4 sample

This program is written in C; it subscribes and gets messages from a topic allowing the use of extended options on the MQSUB call, extending those available on the simpler MQSUB sample: CSQ4BCP2. In addition to the message payload, message properties for each message are received and displayed.

This program takes a variable set of parameters:

- **-t** *Topic string* (required).
- **-o** *Topic object name* (required).
- **-m** *Queue manager name* (optional).
- **-q** *Destination queue name* (optional).
- **-w** *Wait interval on MQGET in seconds* (optional), where *seconds* can have any of the following values:
  - unlimited: MQWI\_UNLIMITED
  - none: No wait
  - *n*: Wait interval in seconds
  - No value specified: When no value is specified, the default is 30 seconds
- **-d** *Subscription name* (optional). Creates or resumes named durable subscription.
- **-k** (optional). Keeps durable subscription on MQCLOSE.

If a queue manager is not specified, CSQ4BCP4 connects to the default queue manager. There is a sample JCL to run the program, it resides in CSQ4BCPE.

### Running the CSQ4BVP1 sample

This program is written in COBOL, it publishes messages to a topic. Start one of the subscriber samples before running this program.

This program takes no parameters. **SYSIN DD** provides the input topic name, queue manager name, and message content.

If a queue manager is not specified, CSQ4BVP1 connects to the default queue manager. There is a sample JCL to run the program, it resides in CSQ4BVPP.

### Running the CSQ4BVP2 sample

This program is written in COBOL, it subscribes to a topic and prints the messages received.

This program takes no parameters. **SYSIN DD** provides the input for topic name and queue manager name.

If a queue manager is not specified, CSQ4BVP1 connects to the default queue manager. There is a sample JCL to run the program, it resides in CSQ4BVPP.

### The Set and Inquire message property sample on z/OS:

The message property sample programs demonstrate the addition of user-defined properties to a message handle, and the inquisition of the properties associated with that message.

The applications use these MQI calls:

- MQCONN
- MQOPEN
- MQPUT
- MQGET
- MQCLOSE
- MQDISC
- MQCRTMH
- MQDLTMH
- MQINQMP
- MQSETMP

The programs are delivered in the C language. The applications run in the batch environment. See Other samples for the batch applications.

The CSQ4BCM1 program is used to inquire the properties of a message handle from a message queue, and it is an example of the use of the MQINQMP API call. The sample gets one message from a queue and then prints all the message handle properties.

The CSQ4BCM2 program is used to set the properties of a message handle on a message queue, and it is an example of the use of the MQSETMP API call. The sample creates a message handle and puts it into the `MsgHandle` field of the `MQGMO` structure. It then puts the message to a queue.

Other examples of inquiring and printing message properties are included in the CSQ4BCG1 and CSQ4BCP4 sample programs.

This topic also provides information on running the Set and Inquire message property samples under the following headings:

- “Running the CSQ4BCM1 sample”
- “Running the CSQ4BCM2 sample”

### Running the CSQ4BCM1 sample

This program takes up to four parameters:

1. The name of the target queue (required).
2. The name of the queue manager (optional).
3. Open options (optional).
4. Close options (optional).

### Running the CSQ4BCM2 sample

This program takes up to six parameters:

1. The name of the target queue (required).
2. The name of the queue manager (optional).
3. Open options (optional).
4. Close options (optional).
5. The name of the target queue manager (optional).
6. The name of the dynamic queue (optional).

The property names, values, and message content are provided through the standard input ( **SYSD DD** ). There is a sample JCL to run the program, it resides in CSQ4BCMP.

---

## Developing .NET applications

IBM MQ classes for .NET allow a program written in the .NET programming framework to connect to IBM MQ as an IBM MQ MQI client or to connect directly to an IBM MQ server.

If you have applications which use Microsoft .NET Framework and want to take advantage of the facilities of IBM MQ, you must use IBM MQ classes for .NET.

The object-oriented IBM MQ .NET interface is different from the MQI interface in that it uses methods of objects rather than using the MQI verbs.

The procedural IBM MQ application programming interface is built around verbs such as those in the following list:

MQCONN, MQDISC, MQOPEN, MQCLOSE,  
MQINQ, MQSET, MQGET, MQPUT, MQSUB

These verbs all take, as a parameter, a handle to the IBM MQ object on which they are to operate. Because .NET is object-oriented, the .NET programming interface turns this round. Your program consists of a set of IBM MQ objects, which you act upon by calling methods on those objects. You can write programs in any language supported by .NET.

When you use the procedural interface, you disconnect from a queue manager by using the call MQDISC( *Hconn*, *CompCode*, *Reason*), where *Hconn* is a handle to the queue manager.

In the .NET interface, the queue manager is represented by an object of class MQQueueManager. You disconnect from the queue manager by calling the Disconnect() method on that class.

```
// declare an object of type queue manager
MQQueueManager queueManager=new MQQueueManager();
...
// do something...
...
// disconnect from the queue manager
queueManager.Disconnect();
```

IBM MQ classes for .NET is a set of classes that enable .NET applications to interact with IBM MQ. They represent the various components of IBM MQ which your application uses, such as queue managers, queues, channels and messages. For details of these classes, see *The IBM MQ .NET classes and interfaces*.

Before you can compile any applications that you write, you must have a .NET Framework installed. For instructions on installing the IBM MQ classes for .NET and the .NET Framework, see “Installing IBM MQ classes for .NET” on page 632.

**Related concepts:**

“Options for connecting to a queue manager” on page 632

There are three modes of connecting IBM MQ classes for .NET to a queue manager. Consider which type of connection best suits your requirements.

“Troubleshooting IBM MQ.NET problems” on page 635

You can use the .NET sample applications to help with troubleshooting problems.

“Writing and deploying IBM MQ .NET programs” on page 647

To use IBM MQ classes for .NET to access IBM MQ queues, you write programs in any language supported by .NET containing calls that put messages onto, and get messages from, IBM MQ queues.

“Developing Microsoft Windows Communication Foundation (WCF) applications with IBM MQ” on page 1178

The Microsoft Windows Communication Foundation (WCF) custom channel for IBM MQ sends and receives messages between WCF clients and services.

“Developing applications for IBM MQ” on page 1

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM MQ support applications written in procedural languages, and object oriented languages and frameworks.

**Related information:**

Technical overview

## Getting started with IBM MQ classes for .NET

IBM MQ classes for .NET allow a program written in the .NET programming framework to connect to IBM MQ as an IBM MQ MQI client or to connect directly to an IBM MQ server.

## Options for connecting to a queue manager

There are three modes of connecting IBM MQ classes for .NET to a queue manager. Consider which type of connection best suits your requirements.

### Client bindings connection

To use IBM MQ classes for .NET as an IBM MQ MQI client, you can install it, with the IBM MQ MQI client, either on the IBM MQ server machine, or on a separate machine. A client bindings connection can use XA or non-XA transactions

### Server bindings connection

When used in server bindings mode, IBM MQ classes for .NET use the queue manager API, rather than communicating through a network. This provides better performance for IBM MQ applications than using network connections.

To use the bindings connection, you must install IBM MQ classes for .NET on the IBM MQ server.

### Managed client connection

A connection made in this mode connects as an IBM MQ client to an IBM MQ server running either on the local or a remote machine.

The IBM MQ classes for .NET connecting in this mode remain in .NET managed code and make no calls to native services. For more information about managed code, refer to Microsoft documentation.

There are a number of limitations to using the managed client. For more information about these, see “Managed client connections” on page 647.

## Installing IBM MQ classes for .NET

IBM MQ classes for .NET, including samples, is installed with IBM MQ. There is a prerequisite of Microsoft.NET Framework.

The latest version of IBM MQ classes for .NET is installed by default as part of the standard IBM MQ installation in the *Java and .NET Messaging and Web Services* feature. For installation instructions, see *Installing IBM MQ server on Windows* or *Installing an IBM MQ client on Windows* systems.

In a multiple installation environment, if you have previously installed the IBM MQ classes for .NET as a support pack, you cannot install IBM MQ unless you first uninstall the support pack. The IBM MQ classes for .NET feature that is installed with IBM MQ contains the same functionality as the support pack.

Sample applications, including source files, are also supplied; see “Sample applications” on page 633.

To run IBM MQ classes for .NET on 32-bit or 64-bit platforms you must have installed Microsoft.NET Framework V3.5 or later.

**Note:** If the Microsoft.NET Framework v3.5 or higher is not installed before installing IBM MQ Version 8.0, then the IBM MQ product installation will continue without error, but the IBM MQ classes for .NET will not be available. If the .NET Framework is installed after installing IBM MQ Version 8.0, then the IBM MQ.NET assemblies must be registered by running the `WMQInstallDir\bin\amqiRegisterdotNet.cmd` script, where `WMQInstallDir` is the directory where IBM MQ Version 8.0 is installed. This script installs the required assemblies in the Global Assembly Cache (GAC). A set of `amqi*.log` files recording the actions taken are created in the `%TEMP%` directory.



For information about using the IBM MQ custom channel for the Microsoft WCF with .NET 3, see "Developing Microsoft Windows Communication Foundation (WCF) applications with IBM MQ" on page 1178

## Sample applications

To run your own .NET applications, use the instructions for the verification programs, substituting your application name in place of the sample applications.

Five sample applications are supplied:

- A put message application
- A get message application
- A 'hello world' application
- A publish/subscribe application
- An application using message properties

All these sample applications are supplied in the C# language, and some are also supplied in C++ and Visual Basic. You can write applications in any language supported by .NET.

### "Put message" program SPUT (nmqsput.cs, mmqsput.cpp, vmqsput.vb)

This program shows how to put a message to a named queue. The program has three parameters:

- The name of a queue (required), for example, SYSTEM.DEFAULT.LOCAL.QUEUE
- The name of a queue manager (optional)
- The definition of a channel (optional), for example, SYSTEM.DEF.SVRCONN/TCP/hostname(1414)

If no queue manager name is given, the queue manager defaults to the default local queue manager. If a channel is defined, it has the same format as the MQSERVER environment variable.

### "Get message" program SGET (nmqsget.cs, mmqsget.cpp, vmqsget.vb)

This program shows how to get a message from a named queue. The program has three parameters:

- The name of a queue (required), for example, SYSTEM.DEFAULT.LOCAL.QUEUE
- The name of a queue manager (optional)
- The definition of a channel (optional), for example, SYSTEM.DEF.SVRCONN/TCP/hostname(1414)

If no queue manager name is given, the queue manager defaults to the default local queue manager. If a channel is defined, it has the same format as the MQSERVER environment variable.

### "Hello World" program (nmqwrlld.cs, mmqwrlld.cpp, vmqwrlld.vb)

This program shows how to put and get a message. The program has three parameters:

- The name of a queue (optional), for example, SYSTEM.DEFAULT.LOCAL.QUEUE or SYSTEM.DEFAULT.MODEL.QUEUE
- The name of a queue manager (optional)
- A channel definition (optional), for example, SYSTEM.DEF.SVRCONN/TCP/hostname(1414)

If no queue name is given, the name defaults to SYSTEM.DEFAULT.LOCAL.QUEUE. If no queue manager name is given, the queue manager defaults to the default local queue manager.

### "Publish/subscribe" program (MQPubSubSample.cs)

This program shows how to use IBM MQ publish/subscribe. It is supplied in C# only. The program has two parameters:

- The name of a queue manager (optional)
- A channel definition (optional)

## "Message properties" program (MQMessagePropertiesSample.cs)

This program shows how to use message properties. It is supplied in C# only. The program has two parameters:

- The name of a queue manager (optional)
- A channel definition (optional)

You can verify your installation by compiling and running these applications.

The sample applications are installed to the following locations, according to the language in which they are written. *MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

### C#

```
MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\nmqswrld.cs
MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\nmqspud.cs
MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\nmqsget.cs
MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\MQPubSubSample.cs
MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\MQMessagePropertiesSample.cs
```

### Managed C++

```
MQ_INSTALLATION_PATH\Tools\dotnet\samples\mcp\mmqswrld.cpp
MQ_INSTALLATION_PATH\Tools\dotnet\samples\mcp\mmqspud.cpp
MQ_INSTALLATION_PATH\Tools\dotnet\samples\mcp\mmqsget.cpp
```

### Visual Basic

```
MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\vmqswrld.vb
MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\vmqspud.vb
MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\vmqsget.vb
MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\xmqswrld.vb
MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\xmqspud.vb
MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\xmqsget.vb
```

To build the sample applications a batch file has been supplied for each language.

### C#

```
MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\bldcssamp.bat
```

The *bldcssamp.bat* file contains a line for each sample, which is all that is necessary to build this sample program:

```
csc /t:exe /r:System.dll /r:amqmdnet.dll /lib:MQ_INSTALLATION_PATH\bin
/out:nmqwrld.exe nmqwrld.cs
```

### Managed C++

```
MQ_INSTALLATION_PATH\Tools\dotnet\samples\mcp\bldmcpsamp.bat
```

The *bldmcpsamp.bat* file contains a line for each sample, which is all that is necessary to build this sample program:

```
cl /clr:oldsyntax MQ_INSTALLATION_PATH\bin mmqwrld.cpp
```

If you want to compile these applications on Microsoft Visual Studio 2003/.NET SDKv1.1, replace the compile command:

```
cl /clr:oldsyntax MQ_INSTALLATION_PATH\bin mmqwrld.cpp
```

with  
cl /clr MQ\_INSTALLATION\_PATH\bin mmqwrld.cpp

## Visual Basic

MQ\_INSTALLATION\_PATH\Tools\dotnet\samples\vb\bldvbsamp.bat

The bldvbsamp.bat file contains a line for each sample, which is all that is necessary to build this sample program:

```
vbc /r:System.dll /r: MQ_INSTALLATION_PATH\bin\amqmdnet.dll /out:vmqwrld.exe vmqwrld.vb
```

## Configuring your queue manager to accept TCP/IP client connections

Configure a queue manager to accept incoming connection requests from the clients.

### About this task

This task explains the basic steps for configuring a queue manager to accept TCP/IP client connections. For a production system, you must also consider the security implications when configuring queue managers.

### Procedure

1. Define a server connection channel:

a. Start the queue manager.

b. Define a sample channel called NET.CHANNEL:

```
DEF CHL('NET.CHANNEL') CHLTYPE(SVRCONN) TRPTYPE(TCP) MCAUSER(' ') +
DESCR('Sample channel for IBM MQ classes for .NET')
```

**Important:** This sample is intended for use in a sandbox environment only, as it does not include any consideration of security implications. For a production system, consider using TLS or a security exit. See *Securing* for more information.

2. Start a listener:

```
runmqclsr -t tcp [-m qmnqme] [-p portnum]
```

**Note:** The square brackets indicate optional parameters; *qmnname* is not required for the default queue manager, and the port number *portnum* is not required if you are using the default (1414).

## Troubleshooting IBM MQ.NET problems

You can use the .NET sample applications to help with troubleshooting problems.

### Using the sample applications

If a program does not complete successfully, run one of the .NET sample applications, and follow the advice given in the diagnostic messages. These sample applications are described in “Sample applications” on page 633.

If the problems continue and you need to contact the IBM service team, you might be asked to turn on the trace facility. For information on using the trace facility, see “Tracing IBM MQ .NET programs” on page 679.

### Error messages

You might see the following common error message:

#### An unhandled exception of type System.IO.FileNotFoundException occurred in unknown module

If this error occurs for either amqmdnet.dll or amqmdxcs.dll, either ensure that both are registered in the Global Assembly Cache or create a configuration file that points to the

amqmdnet.dll and amqmdxcs.dll assemblies. You can examine and change the contents of the assembly cache using mscorecfg.msc, which is supplied as part of the .NET framework.

If the .NET framework was unavailable when IBM MQ was installed, the classes might not be registered in the global assembly cache. You can manually rerun the registration process using the command

```
amqidnet -c MQ_INSTALLATION_PATH\bin\amqidotn.txt -l logfile.txt
```

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

Information about this installation is written to the specified log file ( *logfile.txt* in this example).

## Distributed transactions in .NET

Distributed transactions or global transactions allows client applications to include several different sources of data on two or more networked systems in one transaction.

In distributed transactions, a transaction manager coordinates and manages the transaction among two or more resource managers.

Transactions can be either single phase or two-phase commit process. The single-phase commit is a process where only one resource manager participates in the transaction and two-phase commit process is where there are more than one resource manager participating in the transaction. In the two-phase commit process, the transaction manager sends a prepare call to check whether all the resource managers are prepared to commit. When it receives the acknowledgment from all the resource managers, the commit call is issued. Else, a rollback on the whole transaction happens. See Transaction management and support for more details. The resource managers should inform the transaction managers of their participation in the transaction. When the resource manager informs the transaction manager of its participation, the resource manager gets callbacks from the transaction manager when the transaction is going to commit or roll back.

IBM MQ .NET classes already supports distributed transactions in unmanaged and server bindings mode connections. In these modes, IBM MQ .NET classes delegates all its calls to C extended transaction client, which manages the transaction processing on behalf of .NET.

IBM MQ.NET classes now support distributed transactions in managed mode where IBM MQ .NET Classes uses System.Transactions namespace for the distributed transactions support. The System.Transactions infrastructure makes transactional programming simple and efficient by supporting the transactions initiated in all the resource managers including IBM MQ. The IBM MQ .NET application can put and get messages using .NET implicit transaction programming or explicit transaction programming model. In implicit transactions, the transaction boundaries are created by the application program that decides when to commit, rollback (for explicit transactions) or complete the transaction. In explicit transactions, you have to explicitly specify whether you want to commit, roll back, and complete the transaction.

IBM MQ.NET uses Microsoft distributed transaction coordinator (MS DTC) as the transaction manager, which coordinates and manages the transaction between multiple resource managers. IBM MQ is used as the resource manager. Note that you cannot use TLS with XA transactions. You must use CCDT. For more information, see Using the extended transactional client with TLS channels.

IBM MQ.NET follows the X/Open Distributed Transaction Processing (DTP) model. The X/Open Distributed Transaction Processing model is a distributed transaction processing model proposed by the Open Group, a vendor consortium. This model is a standard among most of the commercial vendors in the transaction processing and database domains. Most of the commercial transaction management products support the X/DTP model.

## Modes of transaction

- “Distributed transactions in managed mode” on page 638
- Distributed transactions for unmanaged mode

## Coordinating transactions in various scenarios

- A connection might participate in several transactions, but only one transaction is active at any point of time.
- During a transaction, the `MQQueueManager.Disconnect` call is honored. In this case the transaction is asked to roll back.
- During a transaction, the `MQQueue.Close` or `MQTopic.Close` call is honored. In this case transaction is asked to roll back.
- The transaction boundaries are created by the application program that decides when to commit, rollback (for explicit transactions) or complete (for implicit transactions) the transaction.
- If the client application breaks during a transaction with an unexpected error before issuing a Put or Get call on a queue or topic call, the transaction is rolled back and a `MQException` is thrown.
- If `MQCC_FAILED` reason code is returned during a Put or Get call on a queue or Topic call, an `MQException` is thrown with reason code and the transaction is rolled. If a prepare call has been already issued by the transaction manager, then IBM MQ .NET returns the prepare request by forcibly rolling back the transaction. Then the transaction manager DTC causes a rollback on current work with all the resource managers in current ambient transactions.
- During a transaction involving multiple resource managers if some environmental reason causes the Put or Get call to hang indefinitely, the transaction manager waits until a stipulated time. After the time is out, it causes the rollback of all current work with all the resource managers in current ambient transactions. If this indefinite wait happens during the prepare phase, the transaction manager might timeout or issue an in-doubt call on the resource in which case the transaction is rolled back.
- Applications using transactions must Put or Get messages under `SYNC_POINT`. If a message Put or Get call is issued under a transactional context that is not under `SYNC_POINT`, the call fails with `MQRC_UNIT_OF_WORK_NOT_STARTED` reason code.

## Behavioral differences between Managed and Unmanaged Client transaction support using Microsoft.NET System.Transactions namespace

Nested Transactions have a `TransactionScope` inside another `TransactionScope`

- IBM MQ .NET fully managed client does support nested `TransactionScope`
- IBM MQ .NET unmanaged client does not support nested `TransactionScope`

Dependent Transactions from `System.Transactions`

- IBM MQ .NET fully managed client does support the dependent transactions facility provided by `System.Transactions`.
- IBM MQ .NET unmanaged client does not support the dependent transactions facility provided by `System.Transactions`.

## Product Samples

New product samples `SimpleXAPut`, and `SimpleXAGet` are available under `WebSphere MQ\tools\dotnet\samples\cs\base`. The samples are C# applications, which demonstrate using `MQPUT` and `MQGET` under Distributed Transactions using `System.Transactions` namespace. For more information about these samples, see “Creating simple put and get messages within a `TransactionScope`” on page 641

## Distributed transactions in managed mode:

IBM MQ .NET classes use System.Transactions namespace for the distributed transactions support in managed mode. In the managed mode, MS DTC coordinates and manages distributed transactions across all the servers enlisted in a transaction.

IBM MQ .NET classes provide an explicit programming model based on the System.Transactions.Transaction class and an implicit programming model using the System.Transactions.TransactionScope, class where the transactions are automatically managed by the infrastructure.

### Implicit Transaction

The following piece of code describes how an IBM MQ .NET application puts a message using .NET implicit transaction programming.

```
Using (TransactionScope scope = new TransactionScope ())
{
 Q.Put (putMsg,pmo);
 scope.Complete ();
}

Q.close();
qMgr.Disconnect();}
```

### Explanation of the code flow of implicit transaction

The code creates *TransactionScope* and puts the message under the scope. It then calls *Complete* to inform the transaction coordinator of the completion of the transaction. The transaction coordinator now issues *prepare* and *commit* to complete the transaction. If an issue is detected, then a *rollback* is called.

### Explicit Transaction

The following code describes how an IBM MQ .NET application puts messages using .NET explicit transaction programming model.

```
MQQueueManager qMgr = new MQQueueManager ("MQQM");
MQQueue Q = QMGR.AccessQueue("Q", MQC.MQOO_OUTPUT+MQC.MQOO_INPUT_SHARED);
MQPutMessageOptions pmo = new MQPutMessageOptions();
pmo.Options = MQC.MQPMO_SYNCPOINT;
MQMessage putMsg1 = new MQMessage();
Using(CommittableTransaction tx = new CommittableTransaction())
{
 try
 {
 Q.Put(MSG,pmo);
 tx.commit();
 }
 catch(Exception)
 {tx.rollback();}
}

Q.close();
qMgr.Disconnect();
}
```

### Explanation of the code flow of explicit transaction

The piece of code creates transaction using *CommittableTransaction* class. It puts a message under that scope and then explicitly calls *commit* to complete the transaction. If there are any issues *rollback* is called.

## Distributed transactions in unmanaged mode:

IBM MQ.NET classes support unmanaged connections (client) using extended transaction client and COM+/MTS as the transaction coordinator, using either implicit or explicit transaction programming model. In the unmanaged mode, IBM MQ .NET classes delegate all its calls to C extended transaction client that manages the transaction processing on behalf of .NET.

The transaction processing is controlled by an external transaction manager, coordinating the global unit of work under the control of the API of the transaction manager. The MQBEGIN, MQCMIT, and MQBACK verbs are unavailable. IBM MQ .NET classes expose this support by way of its unmanaged transport mode (C client). See Configuring XA-compliant transaction managers

MTS is evolved as a transaction processing (TP) system to provide the same features on Windows NT as available in CICS, Tuxedo, and on other platforms. When the MTS is installed, a separate service is added to Windows NT called the Microsoft Distributed Transaction Coordinator (MSDTC). The MSDTC coordinates the transactions that span separate data stores or resources. To work, it requires each data store to implement its own proprietary resource manager.

IBM MQ becomes compatible with MSDTC by implementing an interface (proprietary resource manager interface) where it manages to map DTC XA calls to IBM MQ(X/Open) calls. IBM MQ plays the role of a resource manager.

When a component such as COM+ requests access to an IBM MQ, the COM usually checks with the appropriate MTS context object if a transaction is required. If a transaction is required, the COM informs the DTC and automatically starts an integral IBM MQ transaction for this operation. Then the COM works with the data through the MQMTS software, putting and getting messages as required. The object instance obtained from the COM calls the SetComplete or SetAbort method after all the actions on the data are over. When the application issues SetComplete, the call signals the DTC that the application has completed the transaction and the DTC can go ahead with the two-phase commit process. The DTC then issues calls to MQMTS which in turn issues calls to IBM MQ to commit or roll back the transaction.

## Writing an IBM MQ .NET application using unmanaged client

To run within the context of COM+, a .NET class must inherit from System.EnterpriseServices.ServicedComponent. The rules and recommendations to create assemblies that use serviced components are the following:

**Note:** The following steps are relevant only if you are using System.EnterpriseServices mode.

- The class and method being started in COM+ must both be public (no internal classes, and no protected or static methods).
- The class and method attributes: The TransactionOption attribute dictates the transaction level of the class, that is whether the transactions are disabled, supported, or required. The AutoComplete attribute on the ExecuteUOW() method instructs COM+ to commit the transaction if no unhandled exception is thrown.
- Strong-naming an assembly: The assembly must be strong-named and registered in the Global Assembly Cache (GAC). The assembly is registered in COM+ explicitly or by lazy registration after it is registered in the GAC.
- Registering an assembly in COM+: Prepare the assembly to be exposed to COM clients. Then create a type library by using the Assembly Registration tool, regasm.exe.

```
regasm UnmanagedToManagedXa.dll
```

- Register the assembly into GAC `gacutil /i UnmanagedToManagedXa.dll`.

- Register the assembly in COM+ by using the .NET services installer tool, regsvcs.exe. See the type library created by regasm.exe:

```
Regsvcs /apname:UnmanagedToManagedXa /tlb:UnmanagedToManagedXa.tlb UnmanagedToManagedXa.dll
```

- The assembly is deployed into the GAC, and later it is registered in COM+ by lazy registration. The .NET framework takes care of the registration after the code is run for the first time.

The example code flow using System.EnterpriseServices model and System.Transactions with COM+ are described in the following sections:

### Example code flow using System.EnterpriseServices model

```
using System;
using IBM.WMQ;
using IBM.WMQ.Nmqi;
using System.Transactions;
using System.EnterpriseServices;

namespace UnmanagedToManagedXa
{
 [ComVisible(true)] [System.EnterpriseServices.Transaction(System.EnterpriseServices.TransactionOption.Required)]
 public class MyXa : System.EnterpriseServices.ServicedComponent
 {
 public MQQueueManager QMGR = null;
 public MQQueueManager QMGR1 = null;
 public MQQueue QUEUE = null;
 public MQQueue QUEUE1 = null;
 public MQPutMessageOptions pmo = null;
 public MQMessage MSG = null;

 public MyXa()
 {
 }

 [System.EnterpriseServices.AutoComplete()]
 public void ExecuteUOW()
 {
 QMGR = new MQQueueManager("usemq");

 QUEUE = QMGR.AccessQueue("SYSTEM.DEFAULT.LOCAL.QUEUE",
 MQC.MQOO_INPUT_SHARED +
 MQC.MQOO_OUTPUT +
 MQC.MQOO_BROWSE);

 pmo = new MQPutMessageOptions();
 pmo.Options = MQC.MQPMO_SYNCPOINT;
 MSG = new MQMessage();
 QUEUE.Put(MSG, pmo);
 QMGR.Disconnect();
 }
 }

 public void RunNow()
 {
 MyXa xa = new MyXa();
 xa.ExecuteUOW();
 }
}
```

### Example code flow using System.Transactions for interactions with COM+

```
[STAThread]
public void ExecuteUOW()
{
 Hashtable t1 = new Hashtable();
 t1.Add(MQC.CHANNEL_PROPERTY, "SYSTEM.DEF.SVRCONN");
 t1.Add(MQC.HOST_NAME_PROPERTY, "localhost");
 t1.Add(MQC.PORT_PROPERTY, 1414);
 t1.Add(MQC.TRANSPORT_PROPERTY, MQC.TRANSPORT_MQSERIES_CLIENT);
 TransactionOptions opts = new TransactionOptions();

 using(TransactionScope scope = new TransactionScope(TransactionScopeOption.RequiresNew,
 opts, EnterpriseServicesInteropOption.Full)
 {
 QMGR = new MQQueueManager("usemq", t1);
 }
}
```



```

 QUEUE = QMGR.AccessQueue("SYSTEM.DEFAULT.LOCAL.QUEUE",
 MQC.MQOO_INPUT_SHARED +
 MQC.MQOO_OUTPUT +
 MQC.MQOO_BROWSE);

 pmo = new MQPutMessageOptions();
 pmo.Options = MQC.MQPMO_SYNCPOINT;
 MSG = new MQMessage();
 QUEUE.Put(MSG, pmo);
 scope.Complete();
 }
 QMGR.Disconnect();
}

```

### Creating simple put and get messages within a TransactionScope:

Product sample C# applications are available within IBM MQ. These simple applications demonstrate putting and getting messages within a TransactionScope. At the end of the task, you will be able to put and get messages from a queue or topic.

#### Before you begin

MSDTC service must be running and enabled for XA Transactions.

#### About this task

The example is a simple application, SimpleXAPut and SimpleXAGet. The programs SimpleXAPut and SimpleXAGet are C# applications available within IBM MQ. SimpleXAPut demonstrates using MQPUT, under Distributed Transactions using SystemTransactions namespace. SimpleXAGet demonstrates using MQGET, under Distributed Transactions using SystemTransactions namespace.

SimpleXAPut is located in WebSphere MQ\tools\dotnet\samples\cs\base

#### Procedure

The applications can be run with the command-line parameters from tools\dotnet\samples\cs\base\bin  
 SimpleXAPut.exe -d destinationURI [-h host -p port -l channel -tx transaction -tm mode -n numberOfMsgs]  
 SimpleXAGet.exe -d destinationURI [-h host -p port -l channel -tx transaction -tm mode -n numberOfMsgs]

where the parameters are:

##### **-destinationURI**

This can be queue or topic. For a queue, specify as queue://queueName and for a topic specify as topic://topicName.

##### **-host**

This can be a host name such as localhost or an IP address.

##### **-port**

The port on which the queue manager is running.

##### **-channel**

The connection channel being used. The default is SYSTEM.DEF.SVRCONN

##### **-transaction**

The transaction result, for example commit or rollback.

##### **-mode**

The transport mode, for example managed or unmanaged.

##### **-numberOfMsgs**

The number of messages. The default is 1.

## Example

```
SimpleXAPut -d topic://T01 -h localhost -p 2345 -tx rollback -tm unmanaged
```

```
SimpleXAGet -d queue://Q01 -h localhost -p 2345 -tx rollback -tm unmanaged
```

## Recovering transactions:

This section describes the process of recovering transactions in IBM MQ .NET XA using managed mode.

### Overview

In distributed transaction processing, the transactions can be successfully completed. But, there can be scenarios where a transaction may fail for many reasons. These reasons could include a system failure, hardware failure, network error, incorrect or invalid data, application errors or natural or man-made disasters. It is not possible to prevent transaction failures. The distributed transaction system must be capable of handling these failures. It must be able to detect and correct errors when they occur. This process is known as Transaction Recovery.

An important aspect of the Distributed Transaction Processing is to recover the incomplete or in doubt transactions. It is essential to run the recovery as the Unit of Work part of a particular transaction is held locked until it is recovered. Microsoft.NET from its System.Transactions class library provides the option for recovering incomplete/in-doubt transactions. This recovery support expects Resource Manager to maintain the transaction logs and run the recovery when in need.

### Recovery Model

In the Microsoft .NET transaction recovery model, the Transaction Manager (System.Transactions, or Microsoft Distributed Transaction coordinator (MS DTC), or both), initiates, coordinates, and controls the transaction recovery. The OLE Tx Protocol (the Microsoft XA protocol) based Resource Managers provide the options to configure the DTC to drive, coordinate, and control the recovery for them. To do this, Resource Managers must register XA\_Switch with MS DTC by using native interface.

XA\_Switch provides the entry points of XA functions like xa\_start, xa\_end, and xa\_recover in the Resource Manager to the Distributed Transaction Coordinator.

### Recovery using Microsoft Distributed Transaction coordinator (DTC):

Microsoft Distributed Transaction coordinator provides two kinds of recovery processes.

#### Cold Recovery

Cold recovery is performed if the transaction manager process fails while a connection to a XA resource manager is open. When the transaction manager restarts, it reads the transaction manager logs and re-establishes the connection to the XA resource manager and then initiates recovery.

#### Hot Recovery

Hot recovery is performed if the transaction manager remains up while the connection between the transaction manager and the XA resource manager fails because the XA resource manager or the network fails. After the failure, the transaction manager periodically attempts to reconnect to the XA resource manager. When the connection is re-established, the transaction manager initiates XA recovery.

System.Transactions namespace provides managed implementation of Distributed transactions that are based on MS DTC as the transaction manager. It provides similar features as that of MS DTC's native interface but in fully managed environment. The only difference is about the transaction recovery. System.Transactions expects Resource Managers to drive the recovery by themselves and then coordinate with the Transaction Managers (MS DTC). Resource Manager

must ask for recovery of a particular incomplete transaction and then Transaction Manager accepts it and coordinates based on the actual outcome of that particular transaction.

*IBM MQ .NET Distributed Transaction Processing - Recovering Transaction:*

This section describes how transactions can be recovered with IBM MQ .NET classes.

## Overview

To recover an incomplete transaction the recovery information is required. The transaction recovery information must be logged to storage by the resource managers. IBM MQ .NET classes follow a similar path. The transaction recovery information is logged to a system queue called `SYSTEM.DOTNET.XARECOVERY.QUEUE`.

Transaction recovery in IBM MQ .NET is a two stage process.

1. Logging of transaction recovery information.
  - For every transaction, during the prepare phase a persistent message containing the recovery information is added to `SYSTEM.DOTNET.XARECOVERY.QUEUE`.
  - The message is deleted if the commit call succeeds.
2. Recovering transactions using a monitor application `WmqDotnetXAMonitor`.
  - `WmqDotnetXAMonitor` is a .NET managed application that processes messages in `SYSTEM.DOTNET.XARECOVERY.QUEUE` and recovers incomplete transactions

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.)

## `SYSTEM.DOTNET.XARECOVERY.QUEUE`

This is a system queue that holds transaction recovery information of incomplete transactions. This queue gets created when a queue manager is created.

**Note:** You should not delete `SYSTEM.DOTNET.XARECOVERY.QUEUE` queue.

## `WMQDotnetXAMonitor` Application

IBM MQ .NET XA Monitor application monitors a given queue manager and recovers incomplete transactions if any. Following are deemed as incomplete transactions and are recovered:

### Incomplete Transactions

- If the transaction is prepared but `COMMIT` did not complete within the timeout period.
- If the transaction is prepared but IBM MQ Queue manager has gone down.
- If the transaction is prepared but then Transaction Manager has gone down.

Monitor application must be run from the same system where your IBM MQ .NET client application is running. If there are applications running on multiple systems connecting to the same queue manager, the monitor application must be run from all the systems. Though each client machine has a monitor application running to recover the application, each monitor should be able to identify the message that corresponds to transaction that the current monitor's local MS DTC was coordinating so that it can reenlist and complete it.

Using *WMQDotnetXAMonitor* application:

The *WMQDotnetXAMonitor* application must be run manually. It can be started at any time. You can start it when you see the messages on the *SYSTEM.DOTNET.XARECOVERY.QUEUE* or you can keep it running in the background before you do any transactional work with the applications that are written using IBM MQ .NET classes.

Use the following command to start the monitor application

```
WmqDotnetXAMonitor.exe -m QueueManagerName -n ConnectionName -c ChannelName -h
```

Where

•

**-m *QueueManagerName***

Queue manager name.

Optional

**-n *ConnectionName***

Connection name in host (port) format. Connection Name can contain more than one connection name. Multiple connection names must be given in comma separated list, for example localhost (1414), localhost (1415), localhost (1416). Monitor application run the recovery for each of the connection names specified in the comma separated list.

**-c *ChannelName***

Channel name.

**-h** Heuristic branch completion.

Optional

The monitor application performs the following actions:

1. Checks the queue depth of *SYSTEM.DOTNET.XARECOVERY.QUEUE* at an interval of 100 seconds.
2. If the queue depth is greater than zero the XA monitor browses the queue for messages and checks if the message satisfies the incomplete transaction criteria.
3. If any of the messages satisfy the incomplete transaction criteria, monitor pulls it out, and retrieves the transaction recovery information.
4. It then determines if the recovery information relates to the local MS DTC. If yes then it proceeds to recover the transaction. Otherwise it goes back to browse the next message.
5. It then makes calls to the queue manager to recover the incomplete transaction.

### *WmqDotNETXAMonitor Configuration Settings:*

To monitor the application, inputs can also be provided using the application configuration file. A sample application configuration file is shipped with IBM MQ .NET. This file can be modified according your requirements.

Application Configuration file takes the highest precedence while considering the input values. If input values are provided at both command line and Application Configuration file, then values from application configuration are considered.

Sample application configuration file.

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
<configSections>
<sectionGroup name="IBM.WMQ">
<section name="dnetxa" type="System.Configuration.NameValueFileSectionHandler" />
</sectionGroup>
</configSections>
<IBM.WMQ>
<dnetxa>
<add key="ConnectionName" value="" />
<add key="ChannelName" value="" />
<add key="QueueManagerName" value="" />
<add key="UserId" value="" />
<add key="SecurityExit" value="" />
<add key="SecurityExitUserData" value = "">
</dnetxa>
</dnetxa>
</configuration>
```

### *WmqDotNetXAMonitor Application log:*

Monitor Application creates a log file in the application directory for logging the Monitor's progress and transaction recovery status. Logging starts with the connection name and the channel details to show current queue manager for which the recovery is running.

Once the recovery starts, MessageId of the transaction recovery message, TransactionId of the incomplete transaction and actual outcome of the transaction as per Transaction Manager Coordination will be logged.

Sample log file:

```
Time|ProcessId|ThreadId|WMQ .NET XA Recovery Monitor, Running now for
ConnectionName:xxxx, Time|ProcessId|ThreadId|Channel=xxxx
Time|ProcessId|ThreadId|Current QueueDepth = n
Time|ProcessId|ThreadId|Current MessageId = xxxx
Time|ProcessId|ThreadId|Current Incomplete Transaction being recovered = xxxxx
Time|ProcessId|ThreadId|Actual Outcome of the transaction(as per DTC)= Commit/Roll back
Time|ProcessId|ThreadId|Recovery Completed for TransactionId= xxxxx
Time|ProcessId|ThreadId|Current QueueDepth = n
Time|ProcessId|ThreadId|Current MessageId = xxxx
Time|ProcessId|ThreadId|Current Incomplete Transaction being recovered = xxxxx
Time|ProcessId|ThreadId|Actual Outcome of the transaction(as per DTC)= Commit/Roll back
Time|ProcessId|ThreadId|Recovery Completed for TransactionId= xxxxx
```

## Recovering transactions: Use cases:

There are several different use cases from which transactions might need to be recovered.

- **IBM MQ Application using single DTC and single Queue Manager instance:** In this use case, when you connect to the queue manager and run Unit of Work (UoW) under transaction, and if the transaction fails and becomes incomplete, the monitor application recovers the transaction and completes it.

In this use case, there will be a single instance of monitor application running, as a single Queue manager is associated with the transactions.

- **Multiple IBM MQ Applications using single DTC and single Queue Manager instance:** In this use case, there are more than one WMQ applications under single DTC and all are connecting to the same queue manager and running UoW under transactions.

If the transactions fail and become incomplete, the monitor application recovers them and completes the transactions pertaining to all the applications.

In this use case, a single monitor application runs, as one queue manager is used in transactions.

- **Multiple IBM MQ Applications, multiple DTCs, different Queue Manager instances:** In this use case, there are more than one WMQ applications under different DTCs (that is, each application is running on a different machine) and connecting to different queue managers.

If failure occurs and transaction becomes incomplete, monitor application checks the TransactionManagerWhereabouts in the message to determine the DTC address. If the TransactionManagerWhereabouts value matches with the DTC address under which the monitor is running, it completes the recovery, else, it continues to search until the message corresponding to its DTC is found.

In this use case, there will be only one instance of monitor application running per client (user or computer) as each client has its own queue manager used in transactions.

- **Multiple IBM MQ Applications, multiple DTCs, multiple same Queue Managers instances:** In this use case, there are more than one WMQ applications under different DTCs (that is each application is running on a different machine) and all are connecting to same queue manager.

If failure occurs and transaction becomes incomplete, monitor application verifies the TransactionManagerWhereabouts in the message to check if the DTC address and value match with the DTC under which the monitor is running. If both the values match, it completes the recovery else continues to search until it finds the message corresponding to its DTC.

In this use case, there will be only a single instance of monitor application running per client (user or computer), as each client has its own queue manager association used in transactions.

- **Multiple IBM MQ Applications, single DTC, different Queue Manager instances:** In this use case, there are more than one WMQ applications under a single DTC (that is, on a computer, there are more than one WMQ applications running) and connecting to different queue managers.

If transaction fails and becomes incomplete, monitor application recovers the transaction.

In this use case, there will be as many instances of monitor application running as queue managers connected to, as each application has its own queue manager used in transactions and each of it must be recovered.

**Note:** If the monitor application is not running in the background, you can start it.

## Writing and deploying IBM MQ .NET programs

To use IBM MQ classes for .NET to access IBM MQ queues, you write programs in any language supported by .NET containing calls that put messages onto, and get messages from, IBM MQ queues.

The IBM MQ documentation contains information only on the C#, C++ and Visual Basic languages.

This collection of topics provides information to assist with writing applications to interact with IBM MQ systems. For details of individual classes, see The IBM MQ .NET classes and interfaces.

### Connection differences

The way that you program for IBM MQ.NET has some dependencies on the connection modes that you want to use.

When IBM MQ classes for .NET are used as a managed client, there are a number of differences from a standard IBM MQ MQI client, as some features are not available to a managed client.

IBM MQ.NET determines which connection type to use from the settings that you specify for the connection name, channel name, the customization value `NMQ_MQ_LIB` and the property `MQC.TRANSPORT_PROPERTY`.

#### Managed client connections:

When IBM MQ classes for .NET are used as a managed client, there are a number of differences from a standard IBM MQ MQI client.

The following features are not available to a managed client:

- Channel compression
- Channel exit chaining

If you try to use these features with a managed client, it will return an `MQException`. If the error is detected at the client end of a connection, it will use reason code `MQRC_ENVIRONMENT_ERROR`. If it is detected at the server end, the reason code returned by the server will be used.

Channel exits written for an unmanaged client do not work. You must write new exits specifically for the managed client. Check that there are no invalid channel exits specified in your client channel definition table (CCDT).

The name of a managed channel exit can be up to 999 characters long. However, if you use the CCDT to specify the channel exit name, it is limited to 128 characters.

Communication is supported only over TCP/IP.

When you stop a queue manager using the `endmqm` command, a server-connection channel to a .NET managed client can take longer to close than server-connection channels to other clients.

If you have set `NMQ_MQ_LIB` to managed in order to use managed IBM MQ problem diagnostics, none of the parameters `-i`, `-p`, `-s`, `-b`, or `-c` of the `strmqtrc` command is supported.

A managed .NET application using XA transactions will not work with a z/OS queue manager. A managed .NET client attempting to connect to a z/OS queue manager fails with an error, `MQRC_UOW_ENLISTMENT_ERROR` (`mqr=2354`), on `MQOPEN` call. However, a managed .NET application using XA transactions will work with distributed queue manager.

## Defining which connection type to use:

The connection type is determined by the setting of the connection name, channel name, the customization value `NMQ_MQ_LIB` and the property `MQC.TRANSPORT_PROPERTY`.

You can specify the connection name as follows:

- Explicitly on an `MQueueManager` constructor:

```
public MQueueManager(String queueManagerName, MQLONG Options, string Channel,
string ConnName)
public MQueueManager(String queueManagerName, string Channel, string ConnName)
```
- By setting the properties `MQC.HOST_NAME_PROPERTY` and, optionally, `MQC.PORT_PROPERTY` in a hashtable entry on an `MQueueManager` constructor:

```
public MQueueManager(String queueManagerName, Hashtable properties)
```
- As explicit `MEnvironment` values

```
MEnvironment.Hostname

MEnvironment.Port (optional).
```
- By setting the properties `MQC.HOST_NAME_PROPERTY` and, optionally, `MQC.PORT_PROPERTY` in the `MEnvironment.properties` hashtable.

You can specify the channel name as follows:

- Explicitly on an `MQueueManager` constructor:

```
public MQueueManager(String queueManagerName, MQLONG Options, string Channel,
string ConnName)
public MQueueManager(String queueManagerName, string Channel, string ConnName)
```
- By setting the property `MQC.CHANNEL_PROPERTY` in a hashtable entry on an `MQueueManager` constructor:

```
public MQueueManager(String queueManagerName, Hashtable properties)
```
- As an explicit `MEnvironment` value

```
MEnvironment.Channel
```
- By setting the property `MQC.CHANNEL_PROPERTY` in the `MEnvironment.properties` hashtable.

You can specify the transport property as follows:

- By setting the property `MQC.TRANSPORT_PROPERTY` in a hashtable entry on an `MQueueManager` constructor:

```
public MQueueManager(String queueManagerName, Hashtable properties)
```
- By setting the property `MQC.TRANSPORT_PROPERTY` in the `MEnvironment.properties` hashtable.

Select the connection type you require by using one of the following values:

`MQC.TRANSPORT_MQSERIES_BINDINGS` - connect as server  
`MQC.TRANSPORT_MQSERIES_CLIENT` - connect as non-XA client  
`MQC.TRANSPORT_MQSERIES_XACLIENT` - connect as XA client  
`MQC.TRANSPORT_MQSERIES_MANAGED` - connect as non-XA managed client

You can set the customization value `NMQ_MQ_LIB` to explicitly choose the connection type as shown in the following table.



NMQ_MQ_LIB value	Connection type
mqic.dll	Connect as a non-XA client
mqicxa.dll	Connect as an XA client
mqm.dll	Connect as a server or as a non-XA client
managed	Connect as a non-XA managed client

**Note:** Values of mqic32.dll and mqic32xa.dll are accepted as synonyms of mqic.dll and mqicxa.dll for compatibility with earlier releases. However, mqm.dll and mqm.pdb are only part of the client package from Version 7.1 onwards.

If you choose a connection type which is unavailable in your environment, for example you specify mqic32xa.dll and do not have XA support, IBM MQ.NET throws an exception.

Setting NMQ\_MQ\_LIB to "managed" causes the client to use managed IBM MQ problem diagnostic tests, .NET data conversion, and other managed low-level IBM MQ functions.

All other values for NMQ\_MQ\_LIB cause the .NET process to use unmanaged IBM MQ problem diagnostic tests and data conversion, and other unmanaged low-level IBM MQ functions (assuming an IBM MQ MQI client or server is installed on the system).

IBM MQ.NET chooses the connection type as follows:

1. If MQC.TRANSPORT\_PROPERTY is specified, it connects according to the value of MQC.TRANSPORT\_PROPERTY.

Note, however, that setting MQC.TRANSPORT\_PROPERTY to MQC.TRANSPORT\_MQSERIES\_MANAGED does not guarantee that the client process runs managed. Even with this setting, the client is not managed in the following cases:

- If another thread in the process has connected with MQC.TRANSPORT\_PROPERTY set to something other than MQC.TRANSPORT\_MQSERIES\_MANAGED.
  - If NMQ\_MQ\_LIB is not set to "managed", problem diagnostic tests, data conversion, and other low-level functions are not fully managed (assuming an IBM MQ MQI client or server is installed on the system).
2. If a connection name has been specified without a channel name, or a channel name has been specified without a connection name, it throws an error.
  3. If both a connection name and a channel name have been specified:
    - If NMQ\_MQ\_LIB is set to mqic32xa.dll, it connects as an XA client.
    - If NMQ\_MQ\_LIB is set to managed, it connects as a managed client.
    - Otherwise it connects as a non-XA client.
  4. If NMQ\_MQ\_LIB is specified, it connects according to the value of NMQ\_MQ\_LIB.
  5. If an IBM MQ server is installed, it connects as a server.
  6. If an IBM MQ MQI client is installed, it connects as a non-XA client.
  7. Otherwise, it connects as a managed client.

## Configuration files for IBM MQ classes for .NET

A .NET client application can use an IBM MQ MQI client configuration file and, if you are using the managed connection type, a .NET application configuration file. Settings in the application configuration file have priority.

### Client configuration file

An IBM MQ classes for .NET client application can use a client configuration file in the same way as any other IBM MQ MQI client. This file is typically called mqclient.ini, but you can specify a different file name. For more information about the client configuration file, see [Configuring a client using a configuration file](#).

Only the following attributes in an IBM MQ MQI client configuration file are relevant to IBM MQ classes for .NET. If you specify other attributes, it has no effect.

Stanza	Attribute
CHANNELS	CCSID
CHANNELS	ChannelDefinitionDirectory
CHANNELS	ChannelDefinitionFile
CHANNELS	ServerConnectionParms
ClientExitPath	ExitsDefaultPath
ClientExitPath	ExitsDefaultPath64
MessageBuffer	MaximumSize
MessageBuffer	PurgeTime
MessageBuffer	UpdatePercentage
TCP	ClntRcvBufSize
TCP	ClntSndBufSize
TCP	IPAddressVersion
TCP	KeepAlive

You can override any of these attributes using the appropriate environment variable.

### Application configuration file

If you are running with the managed connection type, you can also override the IBM MQ client configuration file and the equivalent environment variables using the .NET application configuration file.

The .NET application configuration file settings are only acted upon when running with the managed connection type, and are ignored for other connection types.

The .NET application configuration file and its format are defined by Microsoft for general use within the .NET framework, but the particular section names, keys and values mentioned in this documentation are specific to IBM MQ.

The format of the .NET application configuration file is a number of *sections*. Each section contains one or more *keys*, and each key has an associated *value*. The following example shows the sections, keys, and values used in a .NET application configuration file to control the TCP/IP KeepAlive property:

```
<configuration>
 <configSections>
 <section name="TCP" type="System.Configuration.NameValueSectionHandler"/>
 </configSections>
```

```

<TCP>
 <add key="KeepAlive" value="true"></add>
</TCP>
</configuration>

```

The keywords used in the .NET application configuration file section names and keys exactly match the keywords for the Stanzas and Attributes defined in the client configuration file.

See your Microsoft documentation for further information.

### Example code fragment

A C# code fragment demonstrating that an application connects to a queue manager, puts a message on to a queue and receives a reply.

The following C# code fragment demonstrates an application that performs three actions:

1. Connect to a queue manager
2. Put a message onto SYSTEM.DEFAULT.LOCAL.QUEUE
3. Get the message back

It also shows how to change the connection type.

```

// =====
// Licensed Materials - Property of IBM
// 5724-H72
// (c) Copyright IBM Corp. 2003, 2005
// =====
using System;
using System.Collections;

using IBM.WMQ;

class MQSample
{
 // The type of connection to use, this can be:-
 // MQC.TRANSPORT_MQSERIES_BINDINGS for a server connection.
 // MQC.TRANSPORT_MQSERIES_CLIENT for a non-XA client connection
 // MQC.TRANSPORT_MQSERIES_XACLIENT for an XA client connection
 // MQC.TRANSPORT_MQSERIES_MANAGED for a managed client connection
 const String connectionType = MQC.TRANSPORT_MQSERIES_CLIENT;

 // Define the name of the queue manager to use (applies to all connections)
 const String qManager = "your_Q_manager";

 // Define the name of your host connection (applies to client connections only)
 const String hostName = "your_hostname";

 // Define the name of the channel to use (applies to client connections only)
 const String channel = "your_channelname";

 /// <summary>
 /// Initialise the connection properties for the connection type requested
 /// </summary>
 /// <param name="connectionType">One of the MQC.TRANSPORT_MQSERIES_ values</param>
 static Hashtable init(String connectionType)
 {
 Hashtable connectionProperties = new Hashtable();

 // Add the connection type
 connectionProperties.Add(MQC.TRANSPORT_PROPERTY, connectionType);

 // Set up the rest of the connection properties, based on the
 // connection type requested

```

```

switch(connectionType)
{
 case MQC.TRANSPORT_MQSERIES_BINDINGS:
 break;
 case MQC.TRANSPORT_MQSERIES_CLIENT:
 case MQC.TRANSPORT_MQSERIES_XACLIENT:
 case MQC.TRANSPORT_MQSERIES_MANAGED:
 connectionProperties.Add(MQC.HOST_NAME_PROPERTY, hostName);
 connectionProperties.Add(MQC.CHANNEL_PROPERTY, channel);
 break;
}

return connectionProperties;
}
/// <summary>
/// The main entry point for the application.
/// </summary>
[STAThread]
static int Main(string[] args)
{
 try
 {
 Hashtable connectionProperties = init(connectionType);

 // Create a connection to the queue manager using the connection
 // properties just defined
 MQQueueManager qMgr = new MQQueueManager(qManager, connectionProperties);

 // Set up the options on the queue we want to open
 int openOptions = MQC.MQOO_INPUT_AS_Q_DEF | MQC.MQOO_OUTPUT;

 // Now specify the queue that we want to open, and the open options
 MQQueue system_default_local_queue =
 qMgr.AccessQueue("SYSTEM.DEFAULT.LOCAL.QUEUE", openOptions);

 // Define an IBM MQ message, writing some text in UTF format
 MQMessage hello_world = new MQMessage();
 hello_world.WriteUTF("Hello World!");

 // Specify the message options
 MQPutMessageOptions pmo = new MQPutMessageOptions(); // accept the defaults,
 // same as MQPMO_DEFAULT

 // Put the message on the queue
 system_default_local_queue.Put(hello_world, pmo);

 // Get the message back again

 // First define an IBM MQ message buffer to receive the message
 MQMessage retrievedMessage = new MQMessage();
 retrievedMessage.MessageId = hello_world.MessageId;

 // Set the get message options
 MQGetMessageOptions gmo = new MQGetMessageOptions(); //accept the defaults
 //same as MQGMO_DEFAULT

 // Get the message off the queue
 system_default_local_queue.Get(retrievedMessage, gmo);

 // Prove we have the message by displaying the UTF message text
 String msgText = retrievedMessage.ReadUTF();
 Console.WriteLine("The message is: {0}", msgText);

 // Close the queue
 system_default_local_queue.Close();
 }
}

```

```

 // Disconnect from the queue manager
 qMgr.Disconnect();
}

//If an error has occurred,try to identify what went wrong.

//Was it an IBM MQ error?
catch (MQException ex)
{
 Console.WriteLine("An IBM MQ error occurred: {0}", ex.ToString());
}

catch (System.Exception ex)
{
 Console.WriteLine("A System error occurred: {0}", ex.ToString());
}

return 0;
} //end of start
} //end of sample

```

## Setting up the IBM MQ environment

Before you use the client connection to connect to a queue manager, you must set up the IBM MQ environment.

**Note:** This step is not necessary when using IBM MQ classes for .NET in server bindings mode.

The .NET programming interface allows you to use the `NMQ_MQ_LIB` customization value but also includes a class `MQEnvironment`. This class allows you to specify details that are to be used during the connection attempt, such as those in the following list:

- Channel name
- Host name
- Port number
- Channel exits
- SSL parameters
- User ID and password

For full information about the `MQEnvironment` class, see `MQEnvironment.NET` class

To specify the channel name and host name, use the following code:

```

MQEnvironment.Hostname = "host.domain.com";
MQEnvironment.Channel = "client.channel";

```

By default, the clients attempt to connect to a IBM MQ listener at port 1414. To specify a different port, use the code:

```

MQEnvironment.Port = nnnn;

```

## Connecting to and disconnecting from a queue manager

When you have configured the IBM MQ environment, you are ready to connect to a queue manager.

To connect to a queue manager, create a new instance of the `MQQueueManager` class:

```
MQQueueManager queueManager = new MQQueueManager("qMgrName");
```

To disconnect from a queue manager, call the `Disconnect` method on the queue manager:

```
queueManager.Disconnect();
```

You must have `inq` authority on the queue manager when attempting to connect to the queue manager. Without `inq` authority, the connection attempt fails.

If you call the `Disconnect` method, all open queues and processes that you have accessed through that queue manager are closed. However, it is good programming practice to close these resources explicitly when you finish using them. To close the resources, use the `Close` method on the object associated with each resource.

The `Commit` and `Backout` methods on a queue manager replace the `MQCMIT` and `MQBACK` calls that are used with the procedural interface.

## Accessing queues and topics

You can access queues and topics using methods of `MQQueueManager` or appropriate constructors.

To access queues, use the methods of the `MQQueueManager` class. The `MQOD` (object descriptor structure) is collapsed into the parameters of these methods. For example, to open a queue on a queue manager represented by an `MQQueueManager` object called `queueManager`, use the following code:

```
MQQueue queue = queueManager.AccessQueue("qName",
 MQC.MQ00_OUTPUT,
 "qMgrName",
 "dynamicQName",
 "altUserId");
```

The *options* parameter is the same as the `Options` parameter in the `MQOPEN` call.

The `AccessQueue` method returns a new object of class `MQQueue`.

When you have finished using the queue, use the `Close()` method to close it, as in the following example:

```
queue.Close();
```

With IBM MQ .NET, you can also create a queue by using the `MQQueue` constructor. The parameters are exactly the same as for the `accessQueue` method, with the addition of a queue manager parameter specifying the instantiated `MQQueueManager` object to use. For example:

```
MQQueue queue = new MQQueue(queueManager,
 "qName",
 MQC.MQ00_OUTPUT,
 "qMgrName",
 "dynamicQName",
 "altUserId");
```

Constructing a queue object in this way enables you to write your own subclasses of `MQQueue`.

Similarly, you can also access topics using the methods of the `MQQueueManager` class. Use an `AccessTopic()` method to open a topic. This returns a new object of class `MQTopic`. When you have finished using the topic, use the `Close()` method of the `MQTopic` to close it.

You can also create a topic by using an `MQTopic` constructor. There are a number of constructors for topics; for more information see `MQTopic.NET` class.

## Handling messages

Messages are handled using the methods of the queue or topic classes. To build a new message, create a new `MQMessage` object.

Put messages onto queues or topics using the `Put()` method of the `MQueue` or `MTopic` class. Get messages from queues or topics using the `Get()` method of the `MQueue` or `MTopic` class. Unlike the procedural interface, where `MQPUT` and `MQGET` put and get arrays of bytes, the IBM MQ classes for .NET put and get instances of the `MQMessage` class. The `MQMessage` class encapsulates the data buffer that contains the actual message data, together with all the `MQMD` (message descriptor) parameters that describe that message.

To build a new message, create a new instance of the `MQMessage` class and use the `WriteXXX` methods to put data into the message buffer.

When the new message instance is created, all the `MQMD` parameters are automatically set to their default values, as defined in Initial values and language declarations for `MQMD`. The `Put()` method of `MQueue` also takes an instance of the `MQPutMessageOptions` class as a parameter. This class represents the `MQPMO` structure. The following example creates a message and puts it onto a queue:

```
// Build a new message containing my age followed by my name
MQMessage myMessage = new MQMessage();
myMessage.WriteInt(25);

String name = "Charlie Jordan";
myMessage.WriteUTF(name);

// Use the default put message options...
MQPutMessageOptions pmo = new MQPutMessageOptions();

// put the message
!queue.Put(myMessage, pmo);
```

The `Get()` method of `MQueue` returns a new instance of `MQMessage`, which represents the message just taken from the queue. It also takes an instance of the `MQGetMessageOptions` class as a parameter. This class represents the `MQGMO` structure.

You do not need to specify a maximum message size, because the `Get()` method automatically adjusts the size of its internal buffer to fit the incoming message. Use the `ReadXXX` methods of the `MQMessage` class to access the data in the returned message.

The following example shows how to get a message from a queue:

```
// Get a message from the queue
MQMessage theMessage = new MQMessage();
MQGetMessageOptions gmo = new MQGetMessageOptions();
queue.Get(theMessage, gmo); // has default values

// Extract the message data
int age = theMessage.ReadInt();
String name1 = theMessage.ReadUTF();
```

You can alter the number format that the read and write methods use by setting the `encoding` member variable.

You can alter the character set to use for reading and writing strings by setting the `characterSet` member variable.

See `MQMessage.NET` class for more details.

**Note:** The WriteUTF() method of MQMessage automatically encodes the length of the string as well as the Unicode bytes it contains. When your message will be read by another .NET program (using ReadUTF()), this is the simplest way to send string information.

### Handling message properties:

Message properties allow you to select messages, or to retrieve information about a message without accessing its headers. The MQMessage class contains methods to get and set properties.

You can use message properties to allow an application to select messages to process, or to retrieve information about a message without accessing MQMD or MQRFH2 headers. They also facilitate communication between IBM MQ and JMS applications. For more information about message properties in IBM MQ, see Message properties.

The MQMessage class provides a number of methods to get and set properties, according to the data type of the property. The get methods have names of the format Get\*Property, and the set methods have names of the format Set\*Property, where the asterisk (\*) represents one of the following strings:

- Boolean
- Byte
- Bytes
- Double
- Float
- Int
- Int2
- Int4
- Int8
- Long
- Object
- Short
- String

For example, to get the IBM MQ property myproperty (a character string), use the call message.GetStringProperty('myproperty'). You can optionally pass a property descriptor, which IBM MQ will complete.

### Handling errors

Handle errors arising from IBM MQ classes for .NET using try and catch blocks.

Methods in the .NET interface do not return a completion code and reason code. Instead, they throw an exception whenever the completion code and reason code resulting from an IBM MQ call are not both zero. This simplifies the program logic so that you do not have to check the return codes after each call to IBM MQ. You can decide at which points in your program you want to deal with the possibility of failure. At these points, you can surround your code with try and catch blocks, as in the following example:

```
try
{
 myQueue.Put(messageA,PutMessageOptionsA);
 myQueue.Put(messageB,PutMessageOptionsB);
}
catch (MQException ex)
{
 // This block of code is only executed if one of
 // the two put methods gave rise to a non-zero
 // completion code or reason code.
```



```

 Console.WriteLine("An error occurred during the put operation:" +
 "CC = " + ex.CompletionCode +
 "RC = " + ex.ReasonCode);
 Console.WriteLine("Cause exception:" + ex);
}

```

## Getting and setting attribute values

The classes `MQManagedObject`, `MQQueue`, and `MQQueueManager` contain methods that allow you to get and set their attribute values. Note that for `MQQueue`, the methods work only if you specify the appropriate inquire and set flags when you open the queue.

For common attributes, the `MQQueueManager` and `MQQueue` classes inherit from a class called `MQManagedObject`. This class defines the `Inquire()` and `Set()` interfaces.

When you create a new queue manager object by using the *new* operator, it is automatically opened for inquire. When you use the `AccessQueue()` method to access a queue object, that object is not automatically opened for either inquire or set operations, this could cause problems with some types of remote queues. To use the `Inquire` and `Set` methods and to set properties on a queue, you must specify the appropriate inquire and set flags in the `openOptions` parameter of the `AccessQueue()` method.

The inquire and set methods take three parameters:

- selectors array
- intAttrs array
- charAttrs array

You do not need the `SelectorCount`, `IntAttrCount`, and `CharAttrLength` parameters that are found in `MQINQ`, because the length of an array is always known. The following example shows how to make an inquiry on a queue:

```

//inquire on a queue
int [] selectors = new int [2] ;
int [] intAttrs = new int [1] ;
byte [] charAttrs = new byte [MQC.MQ_Q_DESC_LENGTH];
selectors [0] = MQC.MQIA_DEF_PRIORITY;
selectors [1] = MQC.MQCA_Q_DESC;
queue.Inquire(selectors,intAttrs,charAttrs);
ASCIIEncoding enc = new ASCIIEncoding();
String s1 = "";
s1 = enc.GetString(charAttrs);

```

All attributes of these objects can be inquired on. A subset of attributes is exposed as the properties of an object. For a list of object attributes, see `Attributes of objects`. For object properties, see the appropriate class description.

## Multithreaded programs

The .NET runtime environment is inherently multithreaded. IBM MQ classes for .NET allows a queue manager object to be shared across multiple threads but ensures that all access to the target queue manager is synchronized.

Consider a simple program that connects to a queue manager and opens a queue at startup. The program displays a single button on the screen. When a user clicks that button, the program fetches a message from the queue. In this situation, the application initialization occurs in one thread, and the code that executes in response to the button press executes in a separate thread (the user interface thread).

The implementation of IBM MQ .NET ensures that, for a particular connection (MQQueueManager object instance), all access to the target IBM MQ queue manager is synchronized. The default behavior is that a thread that wants to issue a call to a queue manager is blocked until all other calls in progress for that connection are complete. If you require simultaneous access to the same queue manager from multiple threads within your program, create a new MQQueueManager object for each thread that requires concurrent access. (This is equivalent to issuing a separate MQCONN call for each thread.)

If the default connection options are overridden by MQC.MQCNO\_HANDLE\_SHARE\_NONE or MQC.MQCNO\_SHARE\_NO\_BLOCK then the queue manager is no longer synchronized.

## Using a client channel definition table with .NET

You can use a client channel definition table (CCDT) with the .NET classes for IBM MQ. You specify the location of the CCDT in different ways, depending on whether you are using a managed or unmanaged connection.

### Non-XA or XA unmanaged client connection type

With an unmanaged connection type, 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.

### Managed client connection type

With a managed connection type, you can specify the location of the CCDT in three ways:

- Using the .NET application configuration file. In the CHANNELS section, use the keys ChannelDefinitionDirectory to specify the directory where the table is located, and ChannelDefinitionFile to specify the file name.
- 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 in more than one of these ways, the environment variables take priority over the client configuration file, and the .NET Application Configuration File takes priority over both other methods. You can use this feature to specify a standard location in the client configuration file and override it using environment variables or the application configuration file when necessary.

## How a .NET application determines what channel definition to use

In the IBM MQ .NET client environment, the channel definition to be used can be specified in a number of different ways. Multiple specifications of the channel definition can exist. An application derives the channel definition from one or more sources.

If more than one channel definition exists, the one used is selected in the following priority order:

1. Properties specified on the MQQueueManager constructor, either explicitly or by including *MQC.CHANNEL\_PROPERTY* in the properties hashtable
2. A property *MQC.CHANNEL\_PROPERTY* in the MQEnvironment.properties hashtable
3. The property *Channel* in MQEnvironment
4. The .NET application configuration file, section name CHANNELS, key ServerConnectionParms (applies to managed connections only)
5. The *MQSERVER* environment variable
6. The client configuration file, stanza CHANNELS, Attribute ServerConnectionParms
7. The client channel definition table (CCDT). The location of the CCDT is specified in the .NET application configuration file (applies to managed connections only)
8. The client channel definition table (CCDT). The location of the CCDT is specified using the environment variables *MQCHLIB* and *MQCHLTAB*
9. The client channel definition table (CCDT). The location of the CCDT is specified using the client configuration file

For items 1-3, the channel definition is built up field by field from values provided by the application. These values can be provided using different interfaces and multiple values can exist for each one. Field values are added to the channel definition following the priority order given:

1. The value of *connName* on the MQQueueManager constructor
2. Values of properties from the MQQueueManager.properties hashtable
3. Values of properties from the MQEnvironment.properties hashtable
4. Values set as MQEnvironment fields (for example, MQEnvironment.Hostname, MQEnvironment.Port)

For items 4-6, the entire channel definition is supplied as the value. Unspecified fields on the channel definition take the system defaults. No values from other methods of defining channels and their fields are merged with these specifications.

For items 7-9, the entire channel definition is taken from the CCDT. Fields which were not specified explicitly when the channel was defined take the system defaults. No values from other methods of defining channels and their fields are merged with these specifications.

## Using channel exits in IBM MQ .NET

If you use client bindings, you can use channel exits as for any other client connection. If you use managed bindings, you must write an exit program that implements an appropriate interface.

### Client bindings

If you use client bindings, you can use channel exits as described in Channel exits. You cannot use channel exits written for managed bindings.

### Managed bindings

If you use a managed connection, to implement an exit, you define a new .NET class that implements the appropriate interface. Three exit interfaces are defined in the IBM MQ package:

- MQSendExit
- MQReceiveExit
- MQSecurityExit

**Note:** User exits written using these interfaces are not supported as channel exits in the unmanaged environment.

The following sample defines a class that implements all three:

```
class MyMQExits : MQSendExit, MQReceiveExit, MQSecurityExit
{
 // This method comes from the send exit
 byte[] SendExit(MQChannelExit channelExitParms,
 MQChannelDefinition channelDefinition,
 byte[] dataBuffer,
 ref int dataOffset,
 ref int dataLength,
 ref int dataMaxLength)
 {
 // complete the body of the send exit here
 }

 // This method comes from the receive exit
 byte[] ReceiveExit(MQChannelExit channelExitParms,
 MQChannelDefinition channelDefinition,
 byte[] dataBuffer,
 ref int dataOffset,
 ref int dataLength,
 ref int dataMaxLength)
 {
 // complete the body of the receive exit here
 }

 // This method comes from the security exit
 byte[] SecurityExit(MQChannelExit channelExitParms,
 MQChannelDefinition channelDefParms,
 byte[] dataBuffer,
 ref int dataOffset,
 ref int dataLength,
 ref int dataMaxLength)
 {
 // complete the body of the security exit here
 }
}
```

Each exit is passed an MQChannelExit and an MQChannelDefinition object instance. These objects represent the MQCXP and MQCD structures defined in the procedural interface.

The data to be sent by a send exit, and the data received in a security or receive exit is specified using the exit's parameters.

On entry, the data at offset *dataOffset* with length *dataLength* in the byte array *dataBuffer* is the data that is about to be sent by a send exit, and the data received in a security or receive exit. The parameter *dataMaxLength* gives the maximum length (from *dataOffset* ) available to the exit in *dataBuffer*. Note: For a security exit, it is possible for the *dataBuffer* to be null, if this is the first time the exit is called or the partner end elected to send no data.

On return, the value of *dataOffset* and *dataLength* should be set to point to the offset and length within the returned byte array that the .NET classes should then use. For a send exit, this indicates the data that it should send, and for a security or receive exit, the data that should be interpreted. The exit should normally return a byte array; exceptions are a security exit which could elect to send no data, and any exit called with the INIT or TERM reasons. The simplest form of exit that can be written therefore is one which does nothing more than return *dataBuffer*:

The simplest possible exit body is:

```
{
 return dataBuffer;
}
```

### **MQChannelDefinition class**

The userid and password that are specified with the managed .NET client application are set in the IBM MQ .NET MQChannelDefinition class that is passed to the client security exit. The security exit copies the userid and password into the MQCD.RemoteUserIdentifier and MQCD.RemotePassword fields (see "Writing a security exit" on page 355).

### **Specifying channel exits (managed client):**

If you specify a channel name and connection name when creating your MQQueueManager object (either in the MQEnvironment or on the MQQueueManager constructor) you can specify channel exits in two ways.

In order of precedence, these are:

1. Passing hashtable properties MQC.SECURITY\_EXIT\_PROPERTY, MQC.SEND\_EXIT\_PROPERTY or MQC.RECEIVE\_EXIT\_PROPERTY on the MQQueueManager constructor.
2. Setting the MQEnvironment SecurityExit, SendExit or ReceiveExit properties.

If you do not specify a channel name and connection name, the channel exits to use come from the channel definition picked up from a client channel definition table (CCDT). It is not possible to override the values stored in the channel definition. See Client channel definition table and "Using a client channel definition table with .NET" on page 658 for more information about channel definition tables.

In each case, the specification takes the form of a string with the following format:

Assembly\_name(Class\_name)

*Class\_name* is the fully qualified name, including namespace specification, of a .NET class that implements the IBM.WMQ.MQSecurityExit, IBM.WMQ.MQSendExit or IBM.WMQ.MQReceiveExit interface (as appropriate). *Assembly\_name* is the fully qualified location, including file extension, of the assembly that houses the class. The length of the string is limited to 999 characters if you use the properties of MQEnvironment or MQQueueManager. However, if the channel exit name is specified in the CCDT, it is limited to 128 characters. When necessary, the .NET client code loads and creates an

instance of the specified class by parsing the string specification.

### **Specifying channel exit user data (managed client):**

Channel exits can have user data associated with them. If you specify a channel name and connection name when creating your MQQueueManager object (either in the MQEnvironment or on the MQQueueManager constructor) you can specify the user data in two ways.

In order of precedence, these are:

1. Passing hashtable properties MQC.SECURITY\_USERDATA\_PROPERTY, MQC.SEND\_USERDATA\_PROPERTY or MQC.RECEIVE\_USERDATA\_PROPERTY on the MQQueueManager constructor.
2. Setting the MQEnvironment SecurityUserData, SendUserData or ReceiveUserData properties.

If you do not specify a channel name and connection name, the exit user data values to use come from the channel definition picked up from the client channel definition table (CCDT). It is not possible to override the values stored in the channel definition. See Client channel definition table and "Using a client channel definition table with .NET" on page 658 for more information about channel definition tables.

In each case, the specification is a string, limited to 32 characters.

### **Automatic client reconnection in .NET**

You can make your client reconnect automatically to a queue manager during an unexpected connection break.

A client can unexpectedly become disconnected from a queue manager if, for example, the queue manager stops, or the network or server fails.

Without automatic client reconnection, an error is produced when the connection fails. You can use the error code to help you reestablish the connection.

A client that uses the automatic client reconnection facility is called a reconnectable client. To create a reconnectable client, specify certain options called reconnect options while connecting to the queue manager.

If the client application is an IBM MQ .NET client, it can opt to get automatic client reconnection by specifying an appropriate value for CONNECT\_OPTIONS\_PROPERTY when you use the MQQueueManager class to create a queue manager. See Reconnection options for details of CONNECT\_OPTIONS\_PROPERTY values.

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 set of queue managers that are defined with the same QMNAME in the client connection table (See Queue Manager Groups in CCDT for details).

## Transport Layer Security (TLS) support for .NET

IBM MQ classes for .NET client applications support Transport Layer Security (TLS) encryption. The TLS protocol provides communications security over the internet, and allow client/server applications to communicate in a way that is confidential and reliable.

### Related information:

IBM MQ.NET managed client TLS support

Cryptographic security protocols: TLS

### TLS support for the unmanaged .NET client:

TLS support for the unmanaged .NET client is based on the C MQI and GSKit. The C MQI handles the TLS operations and GSKit implements the TLS secure socket protocols.

*Enabling TLS for the unmanaged .NET client:*

TLS is supported only for client connections. To enable TLS, you must specify the CipherSpec to use when communicating with the queue manager, and this must match the CipherSpec set on the target channel.

To enable TLS, specify the CipherSpec using the SSLCipherSpec static member variable of MQEnvironment. The following example attaches to a SVRCONN channel named SECURE.SVRCONN.CHANNEL, which has been set up to require TLS with a CipherSpec of TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA:

```
MQEnvironment.Hostname = "your_hostname";
MQEnvironment.Channel = "SECURE.SVRCONN.CHANNEL";
MQEnvironment.SSLCipherSpec = "TLS_RSA_WITH_AES_128_CBC_SHA";
MQEnvironment.SSLKeyRepository = "C:\mqm\key";
MQQueueManager qmgr = new MQQueueManager("your_Q_manager");
```

See Specifying CipherSpecs for a list of CipherSpecs.

The SSLCipherSpec property can also be set using the MQC.SSL\_CIPHER\_SPEC\_PROPERTY in the hash table of connection properties.

To successfully connect using TLS, the client keystore must be set up with Certificate Authority root certificates chain from which the certificate presented by the queue manager can be authenticated. Similarly, if SSLClientAuth on the SVRCONN channel has been set to MQSSL\_CLIENT\_AUTH\_REQUIRED, the client keystore must contain an identifying personal certificate that is trusted by the queue manager.

*Using the Distinguished Name of the queue manager:*

The queue manager identifies itself using a TLS certificate, which contains a *Distinguished Name* (DN).

An IBM MQ .NET client application can use this DN to ensure that it is communicating with the correct queue manager. A DN pattern is specified using the `sslPeerName` variable of `MQEnvironment`. For example, setting:

```
MQEnvironment.SSLPeerName = "CN=QMGR.*, OU=IBM, OU=WEBSPPHERE";
```

allows the connection to succeed only if the queue manager presents a certificate with a Common Name beginning `QMGR.`, and at least two Organizational Unit names, the first of which must be `IBM` and the second `WEBSPPHERE`.

The `SSLPeerName` property can also be set using the `MQC.SSL_PEER_NAME_PROPERTY` in the hash table of connection properties. For more information about Distinguished Names and rules for setting peer names, refer to *Securing*.

If `SSLPeerName` is set, connections succeed only if it is set to a valid pattern and the queue manager presents a matching certificate.

*Error handling when using TLS:*

The following reason codes can be issued by IBM MQ classes for .NET when connecting to a queue manager using TLS:

**MQRC\_SSL\_NOT\_ALLOWED**

The `SSLCipherSpec` property was set, but `bindings connect` was used. Only client connect supports TLS.

**MQRC\_SSL\_PEER\_NAME\_MISMATCH**

The DN pattern specified in the `SSLPeerName` property did not match the DN presented by the queue manager.

**MQRC\_SSL\_PEER\_NAME\_ERROR**

The DN pattern specified in the `SSLPeerName` property was not valid.

**TLS support for the managed .NET client:**

The managed .NET client uses the Microsoft.NET Framework libraries to implement TLS secure socket protocols. The `Microsoft.System.Net.SecuritySslStream` class operates as a stream over connected TCP sockets and sends and receives data over that socket connection.

The minimum required .NET Framework level is .NET Framework v3.5. The level of Cipher Algorithm support is based on the .NET Framework level that the application is using.

- For applications that are based on .NET Framework level 3.5 and v4.0, the available secure socket protocols are `SSLv3.0` and `TSL v1.0`.
- For applications that are based on .NET Framework level 4.5, the available secure socket protocols are `SSLv3.0`, `TLS v1.1` and `TLSv1.2`.

You might need to move applications that expect higher TLS protocol support to a later version of the framework as defined for Microsoft Security support in the .NET Framework.

The main features of TLS support for the managed .NET client are as follows:

**TLS protocol support**

TLS support for the .NET managed client is defined through the `.NET SSLStream` class, and depends on the .NET Framework that the application is using. For more information see “TLS protocol support for the managed .NET client” on page 666.



**CipherSpec support**

The TLS settings for the .NET managed client are as for the Microsoft.NET TLS steams. For more information see “CipherSpec support for the managed .NET client” on page 666 and “CipherSpec mappings for the managed .NET client” on page 667.

**Key repositories**

The key repository on the client side is a Windows keystore. The server side repository is a Cryptographic Message Syntax (CMS) type of repository. For more information see “Key repositories for the managed .NET client” on page 668.

**Certificates**

You can use self-signed TLS certificates to implement mutual authentication between a client and a queue manager. For more information see “Using certificates for the managed .NET client” on page 669.

**SSLPEERNAME**

In .NET, applications can use the optional SSLPEERNAME attribute to specify a Distinguished Name (DN) pattern. For more information see “SSLPEERNAME” on page 669.

**FIPS compliance**

Enabling FIPS programmatically is not supported by the Microsoft.NET Security library. FIPS enablement is controlled by the Windows Group Policy setting.

**NSA Suite B compliance**

IBM MQ implements RFC 6460. The Microsoft.NET implementation for NSA suite B is 5430. This is supported from .NET Framework 3.5 onwards.

**Secret key reset or renegotiation**

Although the SSLStream class does not support secret key resetting or renegotiation, for consistency with other IBM MQ clients, the .NET managed client allows applications to set SSLKeyResetCount. For more information see “Secret key reset or renegotiation” on page 670.

**Revocation check**

The SSLStream class supports certificate revocation checking, which is automatically done by the certificate chaining engine. For more information see “Revocation check” on page 670.

**IBM MQ security exit support**

The SSLStream class provides limited support for IBM MQ security exits. Querying local and remote certificates to get SSLPeerNamePtr(Subject DN) and SSLRemCertIssNamePtr (Issuer DN) is possible since this is supported in Microsoft.NET. However, there is no support for getting attributes like DNQ, UNSTRUCTUREDNAME and UNSTRUCTUREDADDRESS, so these values cannot be retrieved using the exits.

**Cryptographic hardware support**

Cryptographic hardware is not supported for the managed .NET client.

*TLS protocol support for the managed .NET client:*

IBM MQ.NET TLS support is based on the .NET SSLStream class.

**Note:** TLS protocol support for the managed .NET client depends on the .NET Framework level that the application is using. For more information, see “TLS support for the managed .NET client” on page 664.

For the Microsoft.NET SSLStream class to initialize TLS and perform a hand-shake with the queue manager, one of the required parameters that you must set is **SSLProtocol**, where you must specify the TLS version number, which must be one of the following values:

- SSL3.0
- TLS1.0
- TLS1.2

The value of this parameter is tightly coupled with the Protocol family to which the preferred CipherSpec belongs. When SSLStream starts an TLS handshake with the server (queue manager), it uses the TLS version specified in **SSLProtocol** to identify list of CipherSpecs to be used for negotiation.

IBM MQ.NET does not make any properties available for applications to use to set this value. Instead, IBM MQ uses a mapping table to internally map the CipherSpec set to the Protocol family and identifies the SSLProtocol version to be used. This table shows the mapping each of the supported CipherSpec between Microsoft.NET and IBM MQ, and the Protocol version to which they belong. For more information, see “CipherSpec mappings for the managed .NET client” on page 667.

*CipherSpec support for the managed .NET client:*

The CipherSpec settings for an application are used during the handshake with the server.

IBM MQ clients allow you to set a CipherSpec value that is used during the handshake with the queue manager. Leaving this field blank indicates a plain-text channel without any security on the sockets.

For the IBM MQ.NET managed client, the TLS settings are for the Microsoft.NET SSLStream class. For SSLStream, a CipherSpec, or a preference list of CipherSpecs, can be set only in the Windows group policy, which is a computer-wide setting. SSLStream then uses the specified CipherSpec or preference list during the handshake with the server. In case of other IBM MQ clients, the CipherSpec property can be set in the application on the IBM MQ channel definition and the same setting is used for TLS negotiation. As a result of this restriction, the TLS handshake might negotiate any supported CipherSpec regardless of what is specified in the IBM MQ channel configuration. Therefore, it is likely that this will result in error AMQ9631 on the queue manager. To avoid this error, set the same CipherSpec as the one that you have set in the application as the TLS configuration in the Windows group policy.

The new IBM MQ.NET TLS client code checks only that the correct protocol version was negotiated. The TLS protocol version is derived from the CipherSpec that the application sets and is used for the TLS handshake with the server (queue manager). Hence it is required by design to set the CipherSpec in the IBM MQ.NET managed client application.

## Setting a CipherSpec

There are three ways of setting a CipherSpec:

### MQEnvironment .NET class

The following example shows how to set a CipherSpec with the MQEnvironment class.

```
MQEnvironment.SSLKeyRepository = "*USER";
MQEnvironment.ConnectionName = connectionName;
MQEnvironment.Channel = channelName;
MQEnvironment.properties.Add(MQC.TRANSPORT_PROPERTY, MQC.TRANSPORT_MQSERIES_MANAGED);
MQEnvironment.SSLCipherSpec = "TLS_RSA_WITH_AES_128_CBC_SHA";
```

## TLS CipherSpec property

The following example shows how to set a CipherSpec by adding a hashtable parameter into the MQQueueManager constructor.

```
properties = new Hashtable();
properties.Add(MQC.TRANSPORT_PROPERTY, MQC.TRANSPORT_MQSERIES_MANAGED);
properties.Add(MQC.HOST_NAME_PROPERTY, hostName);
properties.Add(MQC.PORT_PROPERTY, port);
properties.Add(MQC.CHANNEL_PROPERTY, channelName);
properties.Add(MQC.SSL_CERT_STORE_PROPERTY, sslKeyRepository);
properties.Add(MQC.SSL_CIPHER_SPEC_PROPERTY, cipherSpec);
properties.Add(MQC.SSL_PEER_NAME_PROPERTY, sslPeerName);
properties.Add(MQC.SSL_RESET_COUNT_PROPERTY, keyResetCount);
queueManager = new MQQueueManager(queueManagerName, properties);
```

## Windows group policy

When a CipherSpec is set on the Windows group policy, the same CipherSpec must be set for the SSLCipherSpec property value on the SVRCONN channel and in the application. If the Windows group policy is set to the default, that is the group policy is not enabled/edited for CipherSpec setting, applications must set the same default value of the CipherSpec from the Windows group policy TLS configuration in the MQEnvironment class or in the MQQueueManager constructor hashtable properties.

## CCDT usage

IBM MQ.NET only supports Client Channel Definition Tables (.TAB files) that are on a local computer. Existing CCDT files that have a CipherSpec value set can be used for IBM MQ.NET connections. However, the CipherSpec value set on the client connection channel determines the TLS protocol version and also must match the CipherSpec set in the Windows group policy.

### Related concepts:

“Setting up the IBM MQ environment” on page 653

Before you use the client connection to connect to a queue manager, you must set up the IBM MQ environment.

### Related information:

Specifying CipherSpecs

MQEnvironment .NET class

*CipherSpec mappings for the managed .NET client:*

The IBM MQ.NET interface maintains an IBM MQ to Microsoft.NET mapping table that is used to determine the version of the TLS protocol that the managed client needs to use to establish a secure connection with a queue manager.

If a CipherSpec is specified on the SVRCONN channel, then after the TLS handshake is complete, the queue manager tries to match that CipherSpec with the negotiated CipherSpec that the client application is using. If the queue manager cannot find a matching CipherSpec, the communication fails with error AMQ9631.

The IBM MQ.NET interface maintains an IBM MQ to Microsoft.NET CipherSpec mapping table. This table is used to determine the TLS protocol version that client wants to use to establish a secured socket connection with the queue manager. Based on the SSLCipherSpec value, the SSLProtocol version can be TLS v1.0, or TLS v1.2, depending on which version of the Microsoft.NET Framework you are using.

Make sure that you provide the correct the SSLCipherSpec value as specifying an incorrect value might result in failure during the authentication and handshake process. For example, if your application is running on .NET v4.0 you must specify a CipherSpec for TLS v1.0. Specifying a TLS v1.2 CipherSpec might cause failures during authentication.

Table 100. IBM MQ and Microsoft.NET mapping table

IBM MQ CipherSpec	Microsoft.NET CipherSpec	TLS version
TLS_RSA_WITH_AES_128_CBC_SHA	TLS_RSA_WITH_AES_128_CBC_SHA	TLS 1.0
TLS_RSA_WITH_AES_256_CBC_SHA	TLS_RSA_WITH_AES_256_CBC_SHA	TLS 1.0
TLS_RSA_WITH_3DES_EDE_CBC_SHA <sup>1</sup>	TLS_RSA_WITH_3DES_EDE_CBC_SHA <sup>1</sup>	TLS 1.0
TLS_RSA_WITH_AES_128_CBC_SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256	TLS 1.2
TLS_RSA_WITH_AES_256_CBC_SHA256	TLS_RSA_WITH_AES_256_CBC_SHA256	TLS 1.2
ECDHE_RSA_AES_128_CBC_SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256_P256	TLS 1.2
ECDHE_RSA_AES_128_CBC_SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256_P384	TLS 1.2
ECDHE_RSA_AES_128_CBC_SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256_P521	TLS 1.2
ECDHE_RSA_AES_256_CBC_SHA256	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA256_P256	TLS 1.2
ECDHE_RSA_AES_256_CBC_SHA256	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA256_P384	TLS 1.2
ECDHE_RSA_AES_256_CBC_SHA256	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA256_P521	TLS 1.2
ECDHE_ECDSA_AES_128_CBC_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256_P256	TLS 1.2
ECDHE_ECDSA_AES_128_CBC_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256_P384	TLS 1.2
ECDHE_ECDSA_AES_128_CBC_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256_P521	TLS 1.2
ECDHE_ECDSA_AES_256_CBC_SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384_P384	TLS 1.2
ECDHE_ECDSA_AES_256_CBC_SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384_P521	TLS 1.2
ECDHE_ECDSA_AES_128_GCM_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256_P256	TLS 1.2
ECDHE_ECDSA_AES_128_GCM_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256_P384	TLS 1.2
ECDHE_ECDSA_AES_128_GCM_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256_P521	TLS 1.2
ECDHE_ECDSA_AES_256_GCM_SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384_P384	TLS 1.2
ECDHE_ECDSA_AES_256_GCM_SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384_P521	TLS 1.2

**Notes:**

1. This CipherSpec TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA is deprecated. However, it can still be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, you need to either avoid using triple DES, or enable secret key reset when using this CipherSpec.

*Key repositories for the managed .NET client:*

The key repository on the client side of the TLS configuration is a Windows keystore. The server side repository is a Cryptographic Message Syntax (CMS) type of repository.

**Client side**

The client side of the TLS configuration in IBM MQ.NET consists of client side key repository, client certificates, and the choices the application program makes.

- The client side key repository is always a Windows keystore. It can be either a User or a Computer account under which certificates can be stored.
- In the application, you can set either of the following values for the key repository:
  - `"*USER"`: IBM MQ.NET accesses the current user's certificate store to retrieve the client certificates.
  - `"*SYSTEM"`: IBM MQ.NET accesses the local Computer account to retrieve the certificates.
- The client's certificates must be stored in the My certificate store of the User or Computer account. All the server (CA) certificates must be stored in the root directory of the certificate store.

## Server side

The server side key repository is of type CMS and a traceable local computer's directory path must set on the queue manager SSL properties.

*Using certificates for the managed .NET client:*

For client certificates, the IBM MQ managed .NET client accesses the Windows keystore and loads all of the client's certificates that are matched either by certificate label or matched by the string.

When selecting a certificate to use, the IBM MQ managed .NET client always uses the first matching certificate for the SSLStream TLS handshake.

### Matching certificates by certificate label

If you set the certificate label, the IBM MQ managed .NET client searches the Windows certificate store with the given label name to identify the client certificate. It loads all matching certificates and uses the first certificate on the list. There are two options for setting the certificate label:

- The certificate label can set on the MQEnvironment class accessing MQEnvironment.CertificateLabel.
- The certificate label can also be set in a hash table properties, supplied as input parameter with MQQueueManager constructor as shown in the following example.

```
Hashtable properties = new Hashtable();
properties.Add("CertificateLabel", "mycert");
```

The name("CertificateLabel") and the value are case sensitive.

### Matching certificates by string

If certificate label is not set, then the certificate that matches the string "ibmwebspheremq" and the current logged on user (in lower case) is searched for and used.

#### Related information:

MQEnvironment .NET class

Connecting a client to a queue manager securely

#### SSLPEERNAME:

The SSLPEERNAME attribute is used to check the Distinguished Name (DN) of the certificate from the peer queue manager.

In IBM MQ.NET, applications can use SSLPEERNAME to specify a distinguished name pattern as shown in the following example.

```
SSLPEERNAME(CN=QMGR.*, OU=IBM, OU=WEBSHERE)
```

As for other IBM MQ clients, SSLPEERNAME is an optional parameter.

If SSLPEERNAME value is not set, the IBM MQ.NET managed client does not do any Remote(Server) certificate validation and the managed client just accepts the Remote(/server) certificate as-is.

The way in which you set SSLPEERNAME depends on which of the IBM MQ stack offerings you are using.

### IBM MQ classes for .NET

There are three options as follows.

1. Set MQEnvironment.SSLPeerName in the MQEnvironment class.
2. MQEnvironment.properties.Add(MQC.SSL\_PEER\_NAME\_PROPERTY, *value*)

3. Use the queue manager constructor `MQQueueManager` (`String` `queueManagerName`, `Hashtable` `properties`). Supply the `SSLPEERNAME` in the `Hashtable` `properties` as for option 2.

#### **XMS .NET**

Set the SSL peer name in the connection factory:

```
ConnectionFactory.SetStringProperty(XMSC.WMQ_SSL_PEER_NAME, value);
```

**WCF** Include `SslPeerName` as a semicolon separated field in the URI.

#### **Related information:**

`MQEnvironment` .NET class

*Secret key reset or renegotiation:*

The `SSLStream` class does not support secret key reset/renegotiation. However, to be consistent with other IBM MQ clients, the IBM MQ managed .NET client allows applications to set `SSLKeyResetCount`.

When the limit is reached, IBM MQ.NET disconnects from the queue manager and application are notified of that as an exception with `MQRC_CONNECTION_BROKEN` as the reason code. Applications can choose to handle the exception and re-establish connections or enable the `MQCNO_RECONNECT` option for IBM MQ.NET to automatically reconnect to the queue manager.

Enabling the automatic client reconnection facility means that, when the key reset count is reached, all existing connections are brought down and the IBM MQ.NET client re-creates all the connections afresh. For more information about automatic client reconnection, see [Automatic client reconnection](#).

#### **Related information:**

Resetting TLS secret keys

*Revocation check:*

The `SSLStream` class supports certificate revocation checking.

The revocation checking is automatically done by the certificate chaining engine. This applies for both Online Certificate Status Protocol (OCSP) and Certificate Revocation lists (CRLs). The `SSLStream` class uses the certificate revocation that uses only the server specified in the certificate, that is the server is dictated by the certificate itself. It is possible for HTTP CDP extensions and OCSP HTTP requests to proxy through HTTP proxy server.

The way in which you set the revocation check depends on which of the IBM MQ stack offerings you are using.

#### **IBM MQ.NET**

The revocation check can be set by accessing the `MQEnvironment.SSLCertRevocationCheck` property on the `MQEnvironment.cs` class file.

#### **XMS .NET**

The revocation check can be set on the connection factory property context as shown in the following example.

```
ConnectionFactory.SetBooleanProperty(XMSC.WMQ_SSL_CERT_REVOCATION_CHECK, true);
```

**WCF** The revocation check can be set on the URI using the following naming convention.

```
"SslCertRevocationCheck=true"
```

## Configuring TLS for managed IBM MQ .NET:

Configuring TLS for managed IBM MQ .NET consists of creating the signer certificates, then configuring the server side, the client side, and the application program.

### About this task

To configure TLS, you must first create the appropriate signer certificates. Signer certificates can be either self signed or certificates provided by a certificate authority. Although self-signed certificates can be used on a development, test or pre-production system, do not use them on a production system. On a production system, use certificates that you have obtained from a trusted external certificate authority (CA).

### Procedure

1. Create the signer certificates.
  - a. To create self-signed certificates, use one of the following tools provided with IBM MQ : Use either the iKeyman GUI or **iKeycmd** or **runmqakm** from the command line. For more information on using these tools, see Using iKeyman, iKeycmd, runmqakm, and runmqckm.
  - b. To obtain certificates for the queue manager and clients from a certificate authority (CA), follow the instructions in Obtaining personal certificates from a certificate authority.
2. Configure the server side.
  - a. Configure TLS on the queue manager, using GSKit, as described in Connecting a client to a queue manager securely.
  - b. Set the SVRCONN channel TLS attributes:
    - Set **SSLCAUTH** to "REQUIRED/OPTIONAL".
    - Set **SSLCIPH** to an appropriate CipherSpec.

For more information, see "Enabling TLS for the unmanaged .NET client" on page 663.
3. Configure the client side.
  - a. Import the client certificates into the Windows certificate store (under the User/Computer account). IBM MQ .NET accesses client certificates from the Windows certificate store, therefore you must import your certificates into the Windows certificate store to establish a secure socket connection to IBM MQ . For more information about how to access the Windows keystore and import the client side certificates, see Import or export certificates and private keys.
  - b. Supply the CertificateLabel as described in Connecting a client to a queue manager securely.
  - c. If needed, edit the Windows Group Policy to set the CipherSpec, then, for the Windows Group Policy updates to take effect, restart the computer.
4. Configure the application program.
  - a. Set the MQEnvironment or the SSLCipherSpec value to denote the connection as a secured connection. The value that you specify is used to identify the protocol being used (TLS) and must match with any preference that you have specified in the Windows Group Policy.
  - b. Set the SSLKeyRepository property to either "\*SYSTEM" or "\*USER".
  - c. Optional: Set SSLPEERNAME to the distinguished name (DN) of the server certificate.
  - d. Supply the CertificateLabel as described in Connecting a client to a queue manager securely.
  - e. Set any further optional parameters that you require such as KeyResetCount, CertificationRevocationCheck, and enable FIPS.

### Examples of how to set the TLS protocol and TLS key repository

For Base .NET, you can set the TLS protocol and TLS key repository through the MQEnvironment class as shown in the following example:

```
MQEnvironment.SSLCipherSpec = "TLS_RSA_WITH_AES_128_CBC_SHA256";
MQEnvironment.SSLKeyRepository = "*USER";
```

```
MQEnvironment.properties.Add(MQC.SSL_CIPHER_SPEC_PROPERTY, "TLS_RSA_WITH_AES_128_CBC_SHA256")
```

Alternatively, you can set the TLS protocol and TLS key repository by supplying a hashtable as part of the MQQueueManager constructor as shown in the following example.

```
Hashtable properties = new Hashtable();
properties.Add(MQC.SSL_CERT_STORE_PROPERTY, sslKeyRepository);
properties.Add(MQC.SSL_CIPHER_SPEC_PROPERTY, "TLS_RSA_WITH_AES_128_CBC_SHA256")
```

## What to do next

For more information about getting started with developing IBM MQ .NET managed TLS applications, see “Writing a simple application.”

### Related information:

MQEnvironment .NET class

KeyResetCount (MQLONG)

Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows

*Writing a simple application:*

Tips for writing a simple IBM MQ managed .NET TLS application, including examples for setting the SSL properties for connection factories, creating a queue manager instance, connection, session and destination, and sending a test message.

## Before you begin

You must first configure TLS for managed IBM MQ.NET as described in “Configuring TLS for managed IBM MQ .NET” on page 671.

For application program configuration in base .NET, set SSL properties either using the MQEnvironment class or by supplying a hashtable as part of the MQQueueManager constructor.

For application program configuration in XMS .NET, you set the SSL properties on the property context of the connection factories.

## Procedure

1. Set the SSL properties for the connection factories as shown in the following examples.

### Example for IBM MQ.NET

```
properties = new Hashtable();
properties.Add(MQC.TRANSPORT_PROPERTY, MQC.TRANSPORT_MQSERIES_MANAGED);
properties.Add(MQC.HOST_NAME_PROPERTY, hostName);
properties.Add(MQC.PORT_PROPERTY, port);
properties.Add(MQC.CHANNEL_PROPERTY, channelName);
properties.Add(MQC.SSL_CERT_STORE_PROPERTY, sslKeyRepository);
properties.Add(MQC.SSL_CIPHER_SPEC_PROPERTY, cipherSpec);
properties.Add(MQC.SSL_PEER_NAME_PROPERTY, sslPeerName);
properties.Add(MQC.SSL_RESET_COUNT_PROPERTY, keyResetCount);
properties.Add("CertificateLabel", "ibmwebspheremq");
MQEnvironment.SSLCertRevocationCheck = sslCertRevocationCheck;
```

### Example for XMS .NET

```
cf.SetStringProperty(XMSC.WMQ_SSL_KEY_REPOSITORY, "sslKeyRepository");
cf.SetStringProperty(XMSC.WMQ_SSL_CIPHER_SPEC, cipherSpec);
cf.SetStringProperty(XMSC.WMQ_SSL_PEER_NAME, sslPeerName);
cf.SetIntProperty(XMSC.WMQ_SSL_KEY_RESETCOUNT, keyResetCount);
cf.SetBooleanProperty(XMSC.WMQ_SSL_CERT_REVOCATION_CHECK, true);
```



2. Create the queue manager instance, connections, session and destination as shown in the following examples.

#### Example for MQ .NET

```
queueManager = new MQQueueManager(queueManagerName, properties);
Console.WriteLine("done");

// accessing queue
Console.WriteLine("Accessing queue " + queueName + ".. ");
queue = queueManager.AccessQueue(queueName, MQC.MQOO_OUTPUT + MQC.MQOO_FAIL_IF QUIESCING);
Console.WriteLine("done");
```

#### Example for XMS .NET

```
connectionWMQ = cf.CreateConnection();
// Create session
sessionWMQ = connectionWMQ.CreateSession(false, AcknowledgeMode.AutoAcknowledge);

// Create destination
destination = sessionWMQ.CreateQueue(destinationName);

// Create producer
producer = sessionWMQ.CreateProducer(destination);
```

3. Send a message as shown in the following examples.

#### Example for MQ .NET

```
// creating a message object
message = new MQMessage();
message.WriteString(messageString);

// putting messages continuously
for (int i = 1; i <= numberOfMsgs; i++)
{
 Console.WriteLine("Message " + i + " <" + messageString + ">.. ");
 queue.Put(message);
 Console.WriteLine("put");
}
```

#### Example for XMS .NET

```
textMessage = sessionWMQ.CreateTextMessage();
textMessage.Text = simpleMessage;
producer.Send(textMessage);
```

4. Verify the TLS connection. Check the channel status to verify that the TLS connection has been established and is working correctly.

*Configuring trace for SSLStream:*

To capture tracing events and messages relating to the SSLStream class, you must add a configuration section for system diagnostics to the application configuration file for your application.

#### About this task

If you do not add a configuration section for system diagnostics to the application configuration file, the IBM MQ managed .NET client will not capture any events, traces or debugging points relating to TLS and the SSLStream class.

**Note:** Starting IBM MQ tracing using **strmqtrc** does not capture all the required TLS tracing.

#### Procedure

1. Create an application configuration (App.Config) file for your application project.
2. Add a system diagnostics configuration section as shown in the following example.

```

<system.diagnostics>
 <sources>
 <source name="System.Net" tracemode="includehex">
 <listeners>
 <add name="ExternalSourceTrace"/>
 </listeners>
 </source>
 <source name="System.Net.Sockets">
 <listeners>
 <add name="ExternalSourceTrace"/>
 </listeners>
 </source>
 <source name="System.Net.Cache">
 <listeners>
 <add name="ExternalSourceTrace"/>
 </listeners>
 </source>
 <source name="System.Net.Security">
 <listeners>
 <add name="ExternalSourceTrace"/>
 </listeners>
 </source>
 <source name="System.Security">
 <listeners>
 <add name="ExternalSourceTrace"/>
 </listeners>
 </source>
 </sources>
 <switches>
 <add name="System.Net" value="Verbose"/>
 <add name="System.Net.Sockets" value="Verbose"/>
 <add name="System.Net.Cache" value="Verbose"/>
 <add name="System.Security" value="Verbose"/>
 <add name="System.Net.Security" value="Verbose"/>
 </switches>

 <sharedListeners>
 <add name="ExternalSourceTrace" type="IBM.WMQ.ExternalSourceTrace, amqmdnet, Version=8.0.0.0, Culture=
 </sharedListeners>
 <trace autoflush="true"/>
</system.diagnostics>

```

Sample applications for implementing TLS in managed .NET:

Sample applications are provided to show the implementation of TLS for managed .NET in IBM MQ classes for .NET, XMS .NET and IBM MQ custom channel for WCF.

The following table shows the location of the sample applications. *MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

Table 101. Location of sample applications for implementing TLS in managed .NET

IBM MQ.NET stack offering	Location of samples
Base .NET	<i>MQ_INSTALLATION_PATH</i> \Tools\dotnet\samples\cs\base\SimplePut\SimplePut.cs <i>MQ_INSTALLATION_PATH</i> \Tools\dotnet\samples\cs\base\SimpleGet\SimpleGet.cs
XMS .NET	<i>MQ_INSTALLATION_PATH</i> \Tools\dotnet\samples\cs\xms\simple\wmq\SimpleProducer\SimpleProducer.cs <i>MQ_INSTALLATION_PATH</i> \Tools\dotnet\samples\cs\xms\simple\wmq\SimpleConsumer\SimpleConsumer.cs
IBM MQ custom channel for WCF	<i>MQ_INSTALLATION_PATH</i> \Tools\dotnet\samples\cs\wcf\samples\WCF\oneway\service\MQMessagingOneWayService.cs

## Using the .NET Monitor

### Windows

The .NET Monitor is an application similar to an IBM MQ trigger monitor.

**Important:** See Features that can be used only with the primary installation on Windows for important information.

You can create .NET components which are instantiated whenever a message is received on a monitored queue, and which then process that message. The .NET Monitor is started by the **runmqdnm** command and stopped by the **endmqdnm** command. For details of these commands, see **runmqdnm** and **endmqdnm**.

To use the .NET Monitor, you write a component that implements the **IMQObjectTrigger** interface, which is defined in **amqmdnm.dll**.

Components can be either transactional or non-transactional. A transactional component must inherit from **System.EnterpriseServices.ServicedComponent** and be registered as either **RequiresTransaction** or **SupportsTransaction**. It must not be registered as **RequiresNew** as the .NET Monitor has already initiated a transaction.

The component receives **MQQueueManager**, **MQQueue**, and **MQMessage** objects from **runmqdnm**. It may also receive a User Parameter string if one was specified, using the **-u** command-line option, when **runmqdnm** was started. Note that your component receives the contents of a message that arrived on the monitored queue in an **MQMessage** object. It does not have to connect to the queue manager, open the queue, or get the message itself. The component must then process the message as appropriate and return control to the .NET Monitor.

If your component has been written as a transactional component, it registers to commit or roll back the transaction using the facilities provided by **System.EnterpriseServices.ServicedComponent**.

As the component receives **MQQueueManager** and **MQQueue** objects as well as the message, it has complete context information for that message and can, for example, open another queue on the same queue manager without needing to separately connect to IBM MQ.

### Example code fragments: Windows

This topic contains two examples of components which obtain a message from the .NET Monitor and print it, one using transactional processing and the other non-transactional processing. A third example shows common utility routines, applicable to both the first two examples. All the examples are in C#.

#### Example 1: Transactional processing

```
/******
/* Licensed materials, property of IBM */
/* 63H9336 */
/* (C) Copyright IBM Corp. 2005 */
/******
using System;
using System.EnterpriseServices;

using IBM.WMQ;
using IBM.WMQMonitor;

[assembly: ApplicationName("dnmsamp")]

// build:
//
// csc -target:library -reference:amqmdnet.dll;amqmdnm.dll TranAssembly.cs
//
```

```

// run (with dotnet monitor)
//
// runmqdmn -m QMNAME -q QNAME -a dnmsamp.dll -c Tran

namespace dnmsamp
{
 [TransactionAttribute(TransactionOption.Required)]
 public class Tran : ServicedComponent, IMQObjectTrigger
 {
 Util util = null;

 [AutoComplete(true)]
 public void Execute(MQQueueManager qmgr, MQQueue queue,
 MQMessage message, string param)
 {
 util = new Util("Tran");

 if (param != null)
 util.Print("PARAM: '" + param.ToString() + "'");

 util.PrintMessage(message);

 //System.Console.WriteLine("SETTING ABORT");
 //ContextUtil.MyTransactionVote = TransactionVote.Abort;

 System.Console.WriteLine("SETTING COMMIT");
 ContextUtil.SetComplete();
 //ContextUtil.MyTransactionVote = TransactionVote.Commit;
 }
 }
}

```

## Example 2: Non-transactional processing

```

/*****
/* Licensed materials, property of IBM */
/* 63H9336 */
/* (C) Copyright IBM Corp. 2005 */
*****/

using System;

using IBM.WMQ;
using IBM.WMQMonitor;

// build:
//
// csc -target:library -reference:amqmdnet.dll;amqmdnm.dll NonTranAssembly.cs
//
// run (with dotnet monitor)
//
// runmqdmn -m QMNAME -q QNAME -a dnmsamp.dll -c NonTran
namespace dnmsamp
{
 public class NonTran : IMQObjectTrigger
 {
 Util util = null;

 public void Execute(MQQueueManager qmgr, MQQueue queue,
 MQMessage message, string param)
 {
 util = new Util("NonTran");

 try
 {
 util.PrintMessage(message);
 }
 }
 }
}

```

```

 catch (Exception ex)
 {
 System.Console.WriteLine(">>> NonTran\n{0}", ex.ToString());
 }
 }
}

```

### Example 3: Common routines

```

/*****
/* Licensed materials, property of IBM */
/* 63H9336 */
/* (C) Copyright IBM Corp. 2005 */
*****/

using System;

using IBM.WMQ;

namespace dnmsamp
{
 /// <summary>
 /// Summary description for Util.
 /// </summary>
 public class Util
 {
 /* ----- */
 /* Default prefix string of the namespace. */
 /* ----- */
 private string prefixText = "dnmsamp";

 /* ----- */
 /* Constructor that takes the replacement prefix string to use. */
 /* ----- */
 public Util(String text)
 {
 prefixText = text;
 }

 /* ----- */
 /* Display an arbitrary string to the console. */
 /* ----- */
 public void Print(String text)
 {
 System.Console.WriteLine("{0} {1}\n", prefixText, text);
 }

 /* ----- */
 /* Display the content of the message passed to the console. */
 /* ----- */
 public void PrintMessage(MQMessage message)
 {
 if (message.Format.CompareTo(MQC.MQFMT_STRING) == 0)
 {
 try
 {
 string messageText = message.ReadString(message.MessageLength);

 Print(messageText);
 }

 catch(Exception ex)
 {

```

```

 Print(ex.ToString());
 }
}
else
{
 Print("UNRECOGNISED FORMAT");
}
}

/* ----- */
/* Convert the byte array into a hex string. */
/* ----- */
static public string ToHexString(byte[] byteArray)
{
 string hex = "0123456789ABCDEF";

 string retString = "";

 for(int i = 0; i < byteArray.Length; i++)
 {
 int h = (byteArray[i] & 0xF0)>>4;
 int l = (byteArray[i] & 0x0F);

 retString += hex.Substring(h,1) + hex.Substring(l,1);
 }

 return retString;
}
}
}

```

## Compiling IBM MQ .NET programs

Specimen commands to compile .NET applications written in various languages.

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

To build a C# application using IBM MQ classes for .NET, use the following command:

```
csc /t:exe /r:System.dll /r:amqmdnet.dll /lib: MQ_INSTALLATION_PATH\bin /out:MyProg.exe MyProg.cs
```

To build a Visual Basic application using IBM MQ classes for .NET, use the following command:

```
vbc /r:System.dll /r: MQ_INSTALLATION_PATH\bin\amqmdnet.dll /out:MyProg.exe MyProg.vb
```

To build a Managed C++ application using IBM MQ classes for .NET, use the following command:

```
cl /clr MQ_INSTALLATION_PATH\bin Myprog.cpp
```

For other languages, see the documentation supplied by the language vendor.

## Tracing IBM MQ .NET programs

In IBM MQ .NET, you start and control the trace facility as in IBM MQ programs using the MQI.

However, the `-i` and `-p` parameters of the `strmqtrc` command, which allow you to specify process and thread identifiers, and named processes, have no effect.

You normally need to use the trace facility only at the request of IBM service.

See *Using trace on Windows* for information on trace commands.

## Using the stand-alone IBM MQ .NET client

From IBM MQ Version 8.0.0, Fix Pack 2, the IBM MQ .NET client offers you the ability to package and deploy an IBM MQ .NET assembly without needing to use the full IBM MQ client installation on production systems for running your applications.

### About this task

From IBM MQ Version 8.0.0, Fix Pack 2, you can build your IBM MQ .NET applications on a machine where the full IBM MQ client is installed and later package the IBM MQ .NET assembly, that is, `amqmdnet.dll`, along with your application and deploy it on production systems.

The applications that you build and deploy can be the traditional Windows .NET applications, Services, or Microsoft Azure Web/Worker applications.

In such deployments, the IBM MQ .NET client supports only the managed mode of connectivity to a queue manager. The server bindings and unmanaged client mode connectivity are not available as these two modes require a full IBM MQ client installation. Any attempt to use these other two modes results in an application exception.

### Procedure

Referencing the IBM MQ .NET client assembly in applications

- Reference the `amqmdnet.dll` assembly in your application in the same way that you did for earlier releases. Set the **CopyLocal** property of the `amqmdnet` assembly to `True` to ensure that the `amqmdnet` assembly is copied to the `bin` directory of the application. Setting this property also helps the application packaging tool to package the required binary files for deployment on production systems as well as Microsoft Azure PaaS cloud environments.

Adding global transaction support

- Ensure that your application deploys the monitor application `WMQDotnetXAMonitor` on the machine along with the application itself. If an application uses the IBM MQ .NET managed global transaction feature, then it must also deploy the `WMQDotnetXAMonitor` on the machine along with the application itself. This utility is needed for recovering any in-doubt transactions.

Starting and stopping trace

- To start and stop trace, use the application configuration file and an IBM MQ specific trace configuration file.

**Note:** The following steps for generating trace apply to the .NET redistributable managed client as well as the stand-alone .NET client.

You must use the application configuration file and an IBM MQ specific trace configuration file because, since there is no full IBM MQ client installation, the standard tools that are used for starting and stopping trace, `strmqtrc` and `endmqtrc`, are not available.

#### Application configuration file (`app.config` or `web.config`)

Applications need to define the `MQTRACECONFIGFILEPATH` property under the `<appSettings>` section of the application configuration file, that is, the `app.config` or `web.config` file. (The

actual name of the application configuration file depends on the name of your application.) The value of the **MQTRACECONFIGFILEPATH** property specifies the path for the location of the IBM MQ specific trace configuration file, `mqtrace.config`, as show in the following example:

```
<appSettings>
<add key="MQTRACECONFIGFILEPATH" value="C:\MQTRACECONFIG" />
</appSettings>
```

Tracing is disabled if the `mqtrace.config` file is not found in the path that is specified application configuration file. However, First Failure Support Technology™ (FFST) and error logs are created in the application's directory, if the application has authority to write to the current directory.

### IBM MQ specific trace configuration file (`mqtrace.config`)

The `mqtrace.config` file is an XML file that defines properties for starting and stopping trace, the path to the trace files, and the path to the error logs. The following table describes these properties.

Table 102. Properties defined in the `mqtrace.config` file

Attribute	Description
<b>MQTRACELEVEL</b>	0: Stops tracing - this is the default value. 1: Starts tracing with lesser details. 2: Starts tracing with full details - recommended.
<b>MQTRACEPATH</b>	Points to a folder where the trace files will be created. The current directory of the application is used if the path is blank or the <b>MQTRACEPATH</b> attribute is not defined.
<b>MQERRORPATH</b>	Points to a folder where error log files will be created. The current directory of the application is used if the path is blank or the <b>MQERRORPATH</b> attribute is not defined.

The following example shows a sample `mqtrace.config` file:

```
<?xml version="1.0" encoding="utf-8"?>
<traceSettings>
 <MQTRACELEVEL>2</MQTRACELEVEL>
 <MQTRACEPATH>C:\MQTRACEPATH</MQTRACEPATH>
 <MQERRORPATH>C:\MQERRORLOGPATH</MQERRORPATH>
</traceSettings>
```

Tracing can be started and stopped dynamically when an application is running by altering the value of the **MQTRACELEVEL** attribute in the `mqtrace.config` file.

The running application must have create and write permissions for the folder specified by the **MQTRACELEVEL** attribute for generating trace files. Applications that are running in an Microsoft Azure PaaS environment must also ensure similar access permissions since web applications that use an IBM MQ .NET assembly running in Microsoft Azure PaaS might not have create and write permissions. The generation of the trace, first failure data capture (FDC), and error logs fails if the application does not have the required create and write permissions for the specified folder.

### Enabling binding redirection

- To enable compile time binding reference of the IBM MQ .NET assembly to a later version of the assembly, add the `<dependentAssembly>` property to the application configuration file. The following example snippet in the `app.config` file redirects an application that was compiled using the IBM MQ Version 8.0.0, Fix Pack 2 (8.0.0.2) version of the IBM MQ .NET assembly but later a fix pack, IBM MQ Version 8.0.0, Fix Pack 3, was then applied that updated IBM MQ.NET assembly to 8.0.0.3.



```

<runtime>
 <assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1">
 <!-- amqmdnet related binding redirect -->
 <dependentAssembly>
 <assemblyIdentity name="amqmdnet"
 publicKeyToken="dd3cb1c9aae9ec97"
 culture="neutral" />
 <codeBase version="8.0.0.2"
 href="file:///amqmdnet.dll"/>
 <bindingRedirect oldVersion="1.0.0.3-8.0.0.2"
 newVersion="8.0.0.3"/>
 <publisherPolicy apply="no" />
 </dependentAssembly>
 </assemblyBinding>
</runtime>

```

#### Related concepts:

“Using WMQDotnetXAMonitor application” on page 644

The WMQDotnetXAMonitor application must be run manually. It can be started at any time. You can start it when you see the messages on the SYSTEM.DOTNET.XARECOVERY.QUEUE or you can keep it running in the background before you do any transactional work with the applications that are written using IBM MQ .NET classes.

#### Related information:

IBM MQ components and features


Redistributable clients

.NET application runtime - Windows only

---

## Using the Component Object Model Interface (IBM MQ Automation Classes for ActiveX)

The IBM MQ Automation Classes for ActiveX (MQAX) are ActiveX components that provide classes that you can use in your application to access IBM MQ.

 From Version 9.0, IBM MQ support for Microsoft Active X is deprecated. IBM MQ classes for .NET are the recommended replacement technology. For more information, see [Developing .NET applications](#).

MQAX requires an IBM MQ environment and a corresponding IBM MQ application with which to communicate.

It gives your ActiveX application the ability to run transactions and access data on any of your enterprise systems that you can access through IBM MQ.

IBM MQ Automation Classes for ActiveX:

- Give you access to the functions and features of the IBM MQ API, permitting full interconnectivity to other IBM MQ platforms.
- Conform to the normal conventions expected of an ActiveX component.
- Conform to the IBM MQ object model, also available for .NET, C++, Java, and LotusScript®.

MQAX starter samples are provided. You can use these samples initially to check that your installation of the MQAX is successful and that you have the basic IBM MQ environment in place. The samples also demonstrate how MQAX can be used.

## COM and ActiveX scripting

The Component Object Model (COM) is an object-based programming model defined by Microsoft. It specifies how software components can be provided in a way that allows them to locate and communicate with each other irrespective of the computer language in which they are written or their location.

ActiveX is a set of technologies, based on COM, that integrates application development, reusable components, and Internet technologies on the Microsoft Windows platforms. ActiveX components provide interfaces that can be accessed dynamically by applications. An ActiveX scripting client is an application, for example a compiler, that can build or execute a program or script that uses the interfaces provided by ActiveX (or COM) components.

## IBM MQ environment support

IBM MQ Automation Classes for ActiveX can only be used with **32-bit** ActiveX scripting clients.

The COM component can only be used for **32-bit** applications. If you want to write 64-bit COM application, you can use the .NET interface.

To run the MQAX in an IBM MQ server environment you must have Windows 2000 or later installed on your system.

To run the MQAX in an IBM MQ MQI client environment you need IBM MQ MQI client on Windows 2000 or later installed on your system:

The IBM MQ MQI client requires access to at least one IBM MQ server. When both the IBM MQ MQI client and the IBM MQ server are installed on your system, MQAX applications always run against the server. The ActiveX interface to the MQAI is only available in IBM MQ server environments.

## Designing and programming using IBM MQ Automation Classes for ActiveX

### Designing MQAX applications that access non-ActiveX applications

The IBM MQ Automation Classes provide access to the functions of the IBM MQ API. You can therefore benefit from all the advantages that using IBM MQ can bring to your Windows application.

The overall design of your application is the same as for any IBM MQ application, so consider all the design aspects described in the “Design considerations for IBM MQ applications” on page 43 section.

To use the IBM MQ Automation Classes, you code the Windows programs in your application using a language that supports the creation and use of COM objects. For example, Visual Basic, Java, and other ActiveX scripting clients. The classes can then be easily integrated into your application because the IBM MQ objects you need can be coded using the native syntax of the implementation language.

### Using IBM MQ Automation Classes for ActiveX

When designing an ActiveX application that uses IBM MQ Automation Classes for ActiveX, the most important item of information is the message that is sent or received from the remote IBM MQ system. Therefore you must know the format of the items that are inserted into the message. For an MQAX script to a work, both it and the IBM MQ application that picks up or sends the message must know the message structure.

If you are sending a message with an MQAX application and you want to perform data conversion at the MQAX end, you must also know:

- The code page used by the remote system
- The encoding used by the remote system

To keep your code portable, it is good practice to set the code page and encoding, even if they are currently the same in both the sending and receiving systems.

When considering how to structure the implementation of the system you design, remember that your MQAX scripts run on the same machine as the one on which you have either the IBM MQ queue manager or the IBM MQ client installed.

## Programming hints and tips

The following hints and tips are in no significant order. They are subjects that, if relevant to the work you are doing, might save you time.

### Message Descriptor properties

If you manipulate message descriptor properties in a program, it can be better to use the hexadecimal equivalents of the fields.

The information in this section refers to the following properties:

- AccountingToken
- CorrelationId
- GroupId
- MessageId

Where an IBM MQ application is the originator of a message and IBM MQ generates these properties, it is better to use the AccountingTokenHex, CorrelationIdHex, GroupIdHex, and MessageIdHex properties if you want to look at their values, or manipulate them in any way, including passing them back in a message to IBM MQ. The reason for this is that IBM MQ generated values are strings of bytes that have any value from 0 through 255 inclusive, they are not strings of printable characters.

Where your MQAX script is the originator of a message you can use either the AccountingToken, CorrelationId, GroupId, and MessageId properties or their Hex equivalents.

### IBM MQ constants

IBM MQ constants are provided as members of the enum IBM MQ in library MQAX200.

### IBM MQ string constants

IBM MQ string constants are not available when using IBM MQ Automation Classes for ActiveX. You must use the explicit character string for those shown in the following list and any others you might need. The commands must be padded to eight characters using spaces:

*Table 103. IBM MQ string constants and their corresponding character strings.*

String constant	Corresponding character string
MQFMT_NONE	" "
MQFMT_ADMIN	"MQADMIN "
MQFMT_CHANNEL_COMPLETED	"MQCHCOM "
MQFMT_CICS	"MQCICS "
MQFMT_COMMAND_1	"MQCMD1 "
MQFMT_COMMAND_2	"MQCMD2 "
MQFMT_DEAD_LETTER_HEADER	"MQDEAD "

Table 103. IBM MQ string constants and their corresponding character strings. (continued)

String constant	Corresponding character string
MQFMT_DIST_HEADER	"MQHDIST "
MQFMT_EVENT	"MQEVENT "
MQFMT_IMS	"MQIMS "
MQFMT_IMS_VAR_STRING	"MQIMSVS "
MQFMT_MD_EXTENSION	"MQHMDE "
MQFMT_PCF	"MQPCF "
MQFMT_REF_MSG_HEADER	"MQHREF "
MQFMT_RF_HEADER	"MQHRF "
MQFMT_STRING	"MQSTR "
MQFMT_TRIGGER	"MQTRIG "
MQFMT_WORK_INFO_HEADER	"MQHWIH "
MQFMT_XMIT_Q_HEADER	"MQXMIT "

### Null string constants

The IBM MQ constants, used for the initialization of four MQMessage properties, MQMI\_NONE (24 NULL characters), MQCI\_NONE (24 NULL characters), MQGI\_NONE (24 NULL characters), and MQACT\_NONE (32 NULL characters), are not supported by IBM MQ Automation Classes for ActiveX. Setting them to empty strings has the same effect.

For example, to set the various IDs of an MQMessage to these values: *mymessage*. **MessageId** = "" *mymessage*. **CorrelationId** = "" *mymessage*. **AccountingToken** = ""

### Receiving a message from IBM MQ

There are several ways of receiving a message from IBM MQ:

- Polling by issuing a GET followed by a Wait, using the Visual Basic TIMER function.
- Issuing a GET with the Wait option; you specify the wait duration by setting the WaitInterval property. Consider this when, even though you set your system up to run in multithreaded environment, the software running at the time might run only single threaded. This avoids your system locking up indefinitely.

Other threads operate unaffected. However, if your other threads require access to IBM MQ, they require a second connection to IBM MQ using additional MQAX queue manager and queue objects.

Issuing a GET with the Wait option and setting the WaitInterval to MQWI\_UNLIMITED causes your system to lock up until the GET call completes, if the process is single threaded.

## Using data conversion

Two forms of data conversion are supported by IBM MQ Automation Classes for ActiveX - numeric encoding, and character set conversion.

### Numeric encoding

If you set the MQMessage Encoding property, the following methods convert between different numeric encoding systems:

- ReadDecimal2 method
- ReadDecimal4 method
- ReadDouble method
- ReadDouble4 method
- ReadFloat method
- ReadInt2 method
- ReadInt4 method
- ReadLong method
- ReadShort method
- ReadUInt2 method
- WriteDecimal2 method
- WriteDecimal4 method
- WriteDouble method
- WriteDouble4 method
- WriteFloat method
- WriteInt2 method
- WriteInt4 method
- WriteLong method
- WriteShort method
- WriteUInt2 method

The Encoding property can be set and interpreted using the supplied IBM MQ constants. Figure 104 shows an example of these:

```
/* Encodings for Binary Integers */
MQENC_INTEGER_UNDEFINED
MQENC_INTEGER_NORMAL
MQENC_INTEGER_REVERSED

/* Encodings for Decimals */
MQENC_DECIMAL_UNDEFINED
MQENC_DECIMAL_NORMAL
MQENC_DECIMAL_REVERSED

/* Encodings for Floating-Point Numbers */
MQENC_FLOAT_UNDEFINED
MQENC_FLOAT_IEEE_NORMAL
MQENC_FLOAT_IEEE_REVERSED
MQENC_FLOAT_S390
```

*Figure 104. Supplied IBM MQ constants for encoding*

For example, to send an integer from an Intel system to a System/390 operating system in System/390 encoding:

```

Dim msg As New MQMessage 'Define an IBM MQ message for our use..
Print msg. Encoding 'Currently 546 (or X'222')
 'Set the encoding property
 to 785 (or X'311')
msg. Encoding = MQENC_INTEGER_NORMAL OR MQENC_DECIMAL_NORMAL
 OR MQENC_FLOAT_S390
Print msg. Encoding 'Print it to see the change
Dim local_num As long 'Define a long integer
local_num = 1234 'Set it
msg. WriteLong (local_num) 'Write the number into the message

```

## Character set conversion

Character set conversion is necessary when you send a message from one system to another system where the code pages are different. Code page conversion is used by:

- ReadString method
- ReadNullTerminatedString method
- WriteString method
- WriteNullTerminatedString method
- MessageData Property

You must set the MQMessage CharacterSet property to a supported character set value (CCSID).

IBM MQ Automation Classes for ActiveX uses conversion tables to perform character set conversion.

For example, to convert strings automatically to code page 437:

```

Dim msg As New MQMessage 'Define an IBM MQ message
msg.CharacterSet = 437 'Set code page required
msg.WriteString "A character string" 'Put character string in message

```

The WriteString method receives the string data ( A character string in the example) as a Unicode string. It then converts this data from Unicode into code page 437 using the conversion table 34B001B5.TBL.

Characters in the Unicode string that are not supported by code page 437 are given the standard substitution character from code page 437.

Similarly, when you use the ReadString method, the incoming message has a character set established by the IBM MQ Message Descriptor (MQMD) value and there is a conversion from this code page into Unicode before it is passed back to your scripting language.

## Threading

IBM MQ Automation Classes for ActiveX implement a free-threading model where objects can be used between threads.

While MQAX permits the use of MQQueue and MQQueueManager objects, IBM MQ does not currently permit the sharing of handles between different threads.

Attempts to use these on another thread result in an error and IBM MQ returns a return code of MQRC\_HCONN\_ERROR.

**Note:** There is only one MQSession object per process. Using the MQSession CompletionCode and ReasonCode is not recommended in multithreaded environments. The MQSession error values might be overwritten by a second thread between an error being raised and checked on the first thread. Threads are serialized for the duration of each method call or property access. So, issuing a Get with the Wait option causes other threads accessing MQAX objects to be suspended until the operation completes.

## Error handling

This information describes MQAX object properties, how error handling works, rules describing how raising exceptions are handled, and getting a property.

Each MQAX object includes properties to hold error information and a method to reset or clear them. The properties are:

- CompletionCode
- ReasonCode
- ReasonName

The method is:

- ClearErrorCodes

## How error handling works

Your MQAX script or application invokes an MQAX object's method, or accesses or updates a property of the MQAX object:

1. The ReasonCode and CompletionCode in the object concerned are updated.
2. The ReasonCode and CompletionCode in the MQSession object are also updated with the same information.

**Note:** See "Threading" for restrictions on the use of MQSession error codes in threaded applications.

If the CompletionCode is greater than or equal to the ExceptionThreshold property of MQSession, MQAX throws an exception (number 32000). Use this within your script using the On Error (or equivalent) statement to process it.

3. Use the Error function to retrieve the associated error string, which has the form:

```
MQAX: CompletionCode=xxx, ReasonCode=xxx, ReasonName=xxx
```

For more information about how to use the On Error statements, see the documentation for your ActiveX scripting language.

Using the CompletionCode and ReasonCode in the MQSession object is convenient for simple error handlers.

ReasonName property returns the IBM MQ symbolic name for the current value of the ReasonCode.

## Raising exceptions

The following rules describe how raising exceptions are handled:

- Whenever a property or method sets the completion code to a value greater than or equal to the exception threshold (typically set to 2) an exception is raised.
- All method calls and property sets set the completion code.

### Getting a property

This is a special case because the `CompletionCode` and `ReasonCode` are not always updated:

- If a property get succeeds, the object and `MQSession` object `ReasonCode` and `CompletionCode` remain unchanged.
- If a property get fails with a `CompletionCode` of warning, the `ReasonCode` and `CompletionCode` remain unchanged.
- If a property get fails with a `CompletionCode` of error, the `ReasonCode` and `CompletionCode` are updated to reflect the true values, and error processing proceeds as described.

The `MQSession` class has a method `ReasonCodeName` which might be used to replace an IBM MQ reason code with a symbolic name. This is especially useful while developing programs where unexpected errors might occur. However, the name is not ideal for presentation to users.

Each class also has a property `ReasonName`, that returns the symbolic name of the current reason code for that class.

## IBM MQ Automation Classes for ActiveX reference

This section describes the classes of the IBM MQ Automation Classes for ActiveX (MQAX), developed for ActiveX. The classes enable you to write ActiveX applications that can access other applications running in your non-ActiveX environments, using IBM MQ.

### IBM MQ Automation Classes for ActiveX interface

IBM MQ Automation Classes for ActiveX provides predefined numeric ActiveX constants (such as `MQMT_REQUEST`) needed to use the classes.

The ActiveX automation classes consist of the following:

- “MQSession Class” on page 690
- “MQQueueManager class” on page 694
- “MQQueue class” on page 709
- “MQMessage class” on page 728
- “MQPutMessageOptions class” on page 758
- “MQGetMessageOptions class” on page 761
- “MQDistributionList class” on page 763
- “MQDistributionListItem class” on page 769

In addition, IBM MQ Automation Classes for ActiveX provides predefined numeric ActiveX constants (such as `MQMT_REQUEST`) needed to use the classes. These are provided in the enum `MQ` in library `MQAX200`. The constants are a subset of those defined in the IBM MQ C header files (`cmqc*.h`) with some additional IBM MQ Automation Classes for ActiveX Reason codes.



## About IBM MQ Automation Classes for ActiveX classes

Read this information alongside the reference topics under Developing applications reference.

See Features that can be used only with the primary installation on Windows for important information.

The MQSession class provides a root object that contains the status of the last action performed on any of the MQAX objects. See "Error handling" on page 687 for more information.

The MQQueueManager and MQQueue classes provide access to the underlying IBM MQ objects. Methods or property accesses for these classes in general result in calls being made across the IBM MQ MQI.

The MQMessage, MQPutMessageOptions, and MQGetMessageOptions classes encapsulate the MQMD, MQPMO, and MQGMO data structures, and are used to help you send messages to queues and retrieve messages from them.

The MQDistributionList class encapsulates a collection of queues - local, remote, or alias for output. The MQDistributionListItem class encapsulates the MQOR, MQRR, and MQPMR structures and associates them with an owning distribution list.

### Parameter passing

Parameters on method invocations are all passed by value, except where that parameter is an object, in which case it is a reference that is passed.

The class definitions provided list the Data Type for each parameter or property. For many ActiveX clients, such as Visual Basic, if the variable used is not of the required type, the value is automatically converted to or from the required type - providing such a conversion is possible. This follows standard rules of the client; MQAX provides no such conversion.

Many of the methods take fixed-length string parameters, or return a fixed-length character string. The conversion rules are as follows:

- If the user supplies a fixed-length string of the wrong length, as an input parameter or as a return value, the value is truncated or padded with trailing spaces as required.
- If the user supplies a variable-length string of the wrong length as an input parameter, the value is truncated or padded with trailing spaces.
- If the user supplies a variable-length string of the wrong length as a return value, the string is adjusted to the required length (because returning a value destroys the previous value in the string anyway).
- Strings provided as input parameters can contain embedded Nulls.

These classes can be found in the MQAX200 library.

### Object access methods

These methods do not relate directly to any single IBM MQ call. Each of these methods creates an object in which reference information is then held, followed by connecting to or opening an IBM MQ object:

When a connection is made to a queue manager, it holds the 'connection handle' attribute generated by IBM MQ.

When a queue is opened, it holds the 'object handle' attribute generated by IBM MQ.

These IBM MQ attributes are not directly available to the MQAX program.

## Errors

Syntactic errors on parameter passing can be detected at compile time and run time by the ActiveX client. Errors can be trapped using On Error in Visual Basic.

The IBM MQ ActiveX classes all contain two special read-only properties - ReasonCode and CompletionCode. These properties can be read at any time.

An attempt to access any other property, or to issue any method call might generate an error from IBM MQ.

If a property set or method invocation succeeds, the ReasonCode of the owning object is set to MQRC\_NONE and CompletionCode is set to MQCC\_OK.

If the property access or method invocation does not succeed, reason and completion codes are set in these fields.

## MQSession Class

This is the root class for IBM MQ Automation Classes for ActiveX.

There is always only one MQSession object per ActiveX client process. An attempt to create a second object creates a second reference to the original object.

## Creation

New creates a new MQSession object.

## Syntax

```
Dim mqsess As New MQSession Set mqsess = New MQSession
```

## Properties

- “CompletionCode property” on page 691.
- “ExceptionThreshold property” on page 691.
- “ReasonCode property” on page 692.
- “ReasonName property” on page 692.

## Method

- “AccessGetMessageOptions method” on page 693.
- “AccessMessage method” on page 693.
- “AccessPutMessageOptions method” on page 693.
- “AccessQueueManager method” on page 693.
- “ClearErrorCodes method” on page 694.
- “ReasonCodeName method” on page 694.

**CompletionCode property:**

Read-only. Returns the IBM MQ completion code set by the most recent method or property set issued against any IBM MQ object.

It is reset to MQCC\_OK when a method or a property set is invoked successfully against any MQAX object.

An error event handler can inspect this property to diagnose the error, without having to know which object was involved.

Using the CompletionCode and ReasonCode in the MQSession object is very convenient for simple error handlers.

**Note:** See “Threading” on page 687 for restrictions on the use of MQSession error codes in threaded applications.

**Defined in:**

MQSession class

**Data Type:**

Long

**Values:**

- MQCC\_OK
- MQCC\_WARNING
- MQCC\_FAILED

**Syntax:**

To get: *completioncode*& = MQSession.CompletionCode

**ExceptionThreshold property:**

Read-write. Defines the level of IBM MQ error for which MQAX will throw an exception. Defaults to MQCC\_FAILED. A value greater than MQCC\_FAILED effectively prevents exception processing, leaving the programmer to perform checks on the CompletionCode and ReasonCode.

**Defined in:** MQSession class

**Data Type:** Long

**Values:**

- Any, but consider MQCC\_WARNING, MQCC\_FAILED or greater.

**Syntax:**

To get: *ExceptionThreshold*& = MQSession. **ExceptionThreshold**

To set: MQSession. **ExceptionThreshold** = *ExceptionThreshold*\$

**ReasonCode property:**

Read-only. Returns the reason code set by the most recent method or property set issued against any IBM MQ object.

An error event handler can inspect this property to diagnose the error, without having to know which object was involved.

Using the CompletionCode and ReasonCode in the MQSession object is very convenient for simple error handlers.

**Note:** See “Threading” on page 687 for restrictions on the use of MQSession error codes in threaded applications.

**Defined in:** MQSession class

**Data Type:** Long

**Values:**

- See Reason (MQLONG) and the additional MQAX values listed under “Reason codes IBM MQ Automation Classes for ActiveX” on page 778.

**Syntax:** To get: *reasoncode*& = *MQSession*. **ReasonCode**

**ReasonName property:**

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC\_QMGR\_NOT\_AVAILABLE".

**Note:** See “Threading” on page 687 for restrictions on the use of MQSession error codes in threaded applications.

**Defined in:** MQSession class

**Data Type:** String

**Values:**

- See API reason codes.

**Syntax:** To get: *reasonname*\$ = *MQSession* .**ReasonName**

**AccessGetMessageOptions method:**

Creates a new MQGetMessageOptions object.

**Defined in:**

MQSession class

**Syntax:**

*gmo* = *MQSession* .AccessGetMessageOptions()

**AccessMessage method:**

Creates a new MQMessage object.

**Defined in:**

MQSession class

**Syntax:**

*msg* = *MQSession* .AccessMessage()

**AccessPutMessageOptions method:**

Creates a new MQPutMessageOptions object.

**Defined in:**

MQSession class

**Syntax:**

*pmo* = *MQSession* .AccessPutMessageOptions()

**AccessQueueManager method:**

Creates a new MQQueueManager object and connects it to a real queue manager by means of the IBM MQ MQI client or IBM MQ server. As well as performing a connect, this method also performs an open for the queue manager object.

When both the IBM MQ MQI client and IBM MQ server are installed on your system, MQAX applications will run against the server by default. To run MQAX against the client, the client bindings library must be specified in the GMQ\_MQ\_LIB environment variable, for example, set GMQ\_MQ\_LIB=mqic.dll.

For a client only installation, it is not necessary to set the GMQ\_MQ\_LIB environment variable. When this variable is not set, IBM MQ attempts to load amqzst.dll. If this DLL is not present (as is the case in a client only installation), IBM MQ attempts to load mqic.dll.

If successful it sets the MQQueueManager's ConnectionStatus to TRUE.

A queue manager can be connected to at most one MQQueueManager object per ActiveX instance.

If the connection to the queue manager fails, an error event is raised, and the MQSession object's ReasonCode and CompletionCode are set.

**Defined in:** MQSession class

**Syntax:** *set qm* = *MQSession* .AccessQueueManager ( *Name\$* )

**Parameter:** *Name\$* String. Name of Queue Manager to be connected to.

### ClearErrorCodes method:

Resets the CompletionCode to MQCC\_OK and the ReasonCode to MQRC\_NONE.

Defined in: MQSession class

### Syntax:

Call `MQSession.ClearErrorCodes()`

### ReasonCodeName method:

Returns the name of the reason code with the given numeric value. It is useful to give clearer indications of error conditions to users. The name is still somewhat cryptic (for example, ReasonCodeName(2059) is **MQRC\_Q\_MGR\_NOT\_AVAILABLE**), so where possible errors should be caught and replaced with descriptive text appropriate to the application.

Defined in: MQSession class

Syntax: `errname$ = MQSession.ReasonCodeName ( reasonCode& )`

Parameter: `reasoncode&` Long. The reason code for which the symbolic name is required.

## MQQueueManager class

This class represents a connection to a queue manager. The queue manager can be running locally (an IBM MQ server) or remotely with access provided by the IBM MQ client. An application must create an object of this class and connect it to a queue manager. When an object of this class is destroyed it is automatically disconnected from its queue manager.

### Containment

MQQueue class objects are associated with this class.

New creates a new MQQueueManager object and sets all the properties to initial values. Alternatively use the AccessQueueManager method of the MQSession class.

### Creation

New creates a **new** MQQueueManager object and sets all the properties to initial values. Alternatively use the AccessQueueManager method of the MQSession class.

### Syntax

Dim `mgr` As New MQQueueManager set `mgr` = New MQQueueManager

### Properties

- "AlternateUserId property" on page 696.
- "AuthorityEvent property" on page 696.
- "BeginOptions property" on page 697.
- "ChannelAutoDefinition property" on page 697.
- "ChannelAutoDefinitionEvent property" on page 697.
- "ChannelAutoDefinitionExit property" on page 698.
- "CharacterSet property" on page 698.
- "CloseOptions property" on page 698.

- “CommandInputQueueName property” on page 698.
- “CommandLevel property” on page 699.
- “CompletionCode property” on page 699.
- “ConnectionHandle property” on page 699.
- “ConnectionStatus property” on page 699.
- “ConnectOptions property” on page 700.
- “DeadLetterQueueName property” on page 700.
- “DefaultTransmissionQueueName property” on page 700.
- “Description property” on page 700.
- “DistributionLists property” on page 701.
- “InhibitEvent property” on page 701.
- “IsConnected property” on page 701.
- “IsOpen property” on page 702.
- “LocalEvent property” on page 702.
- “MaximumHandles property” on page 702.
- “MaximumMessageLength property” on page 702.
- “MaximumPriority property” on page 703.
- “MaximumUncommittedMessages property” on page 703.
- “Name property” on page 703.
- “ObjectHandle property” on page 703.
- “PerformanceEvent property” on page 704.
- “Platform property” on page 704.
- “ReasonCode property” on page 704.
- “ReasonName property” on page 705.
- “RemoteEvent property” on page 705.
- “StartStopEvent property” on page 705.
- “SyncPointAvailability property” on page 706.
- “TriggerInterval property” on page 706.

## Methods

- “AccessQueue method” on page 706.
- “AddDistributionList method” on page 707.
- “Backout method” on page 707.
- “Begin method” on page 707.
- “ClearErrorCodes method” on page 708.
- “Commit method” on page 708.
- “Connect method” on page 708.
- “Disconnect method” on page 708.

## Property Access

The following properties can be accessed at any time .

- “AlternateUserId property” on page 696.
- “CompletionCode property” on page 699.
- “ConnectionStatus property” on page 699.
- “ReasonCode property” on page 704.

The remaining properties can be accessed only if the object is connected to a queue manager, and the user ID is authorized to inquire against that queue manager. If an alternate user ID is set and the current user ID is authorized to use it, the alternate user ID is checked for authorization for inquire instead

If these conditions do not apply, IBM MQ Automation Classes for ActiveX attempts to connect to the queue manager and open it for inquire automatically. If this is unsuccessful, the call sets a CompletionCode of MQCC\_FAILED and one of the following ReasonCodes:

- MQRC\_CONNECTION\_BROKEN
- MQRC\_NOT\_AUTHORIZED
- MQRC\_Q\_MGR\_NAME\_ERROR
- MQRC\_Q\_MGR\_NOT\_AVAILABLE

**AlternateUserId property:**

Read-write. The alternate user ID to be used to validate access to the queue manager attributes.

This property must not be set if IsConnected is TRUE.

This property cannot be set while the object is open.

**Defined in:** MQQueueManager class

**Data Type:** String of 12 characters

**Syntax:** To get: *altuser\$* = *MQQueueManager* .**AlternateUserId** To set: *MQQueueManager* .**AlternateUserId** = *altuser\$*

**AuthorityEvent property:**

Read-only. The MQI AuthorityEvent attribute.

**Defined in:**  
MQQueueManager class

**Data Type:**  
Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *authevent* = *MQQueueManager* .**AuthorityEvent**



**BeginOptions property:**

Read-write. These are the options that apply to the Begin method. Initially MQBO\_NONE.

**Defined in:**

MQQueueManager class

**Data Type:**

Long

**Values:**

- MQBO\_NONE

**Syntax:** To get: *beginoptions&=MQQueueManager*. **BeginOptions**

To set: *MQQueueManager* **.BeginOptions** = *beginoptions&*

**ChannelAutoDefinition property:**

Read-only. This controls whether automatic channel definition is permitted.

**Defined in:**

MQQueueManager class

**Data Type:**

Long

**Values:**

- MQCHAD\_DISABLED
- MQCHAD\_ENABLED

**Syntax:** To get: *channelautodef&= MQQueueManager*. **ChannelAutoDefinition**

**ChannelAutoDefinitionEvent property:**

Read-only. This controls whether automatic channel definition events are generated.

**Defined in:**

MQQueueManager class

**Data Type:**

Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *channelautodefevent&=MQQueueManager*. **ChannelAutoDefinitionEvent**

**ChannelAutoDefinitionExit property:**

Read-only. The name of the user exit used for automatic channel definition.

**Defined in:**

MQQueueManager class

**Data Type:**

String

**Syntax:** To get: *channelautodefexit\$* = *MQQueueManager*. **ChannelAutoDefinitionExit**

**CharacterSet property:**

Read-only. The MQI CodedCharSetId attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Syntax:** To get: *characterset&* = *MQQueueManager*. **CharacterSet**

**CloseOptions property:**

Read-write. Options used to control what happens when the queue manager is closed. The initial value is MQCO\_NONE.

**Defined in:**

MQQueueManager class

**Data Type:**

Long

**Values:**

- MQCO\_NONE

**Syntax:** To get: *closeopt&* = *MQQueueManager*. **CloseOptions**

To set: *MQQueueManager*. **CloseOptions** = *closeopt&*

**CommandInputQueueName property:**

Read-only. The MQI CommandInputQName attribute.

**Defined in:** MQQueueManager class

**Data Type:** String of 48 characters

**Syntax:** To get: *commandinputqname\$* = *MQQueueManager*. **CommandInputQueueName**

**CommandLevel property:**

Read-only. Returns the version and level of the IBM MQ queue manager implementation (MQI CommandLevel attribute)

**Defined in:** MQQueueManager class

**Data Type:** Long

**Syntax:** To get: *level*& = *MQQueueManager* .**CommandLevel**

**CompletionCode property:**

Read-only. Returns the completion code set by the last method or property access issued against the object.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- MQCC\_OK
- MQCC\_WARNING
- MQCC\_FAILED

**Syntax:** To get: *completioncode*& = *MQQueueManager* .**CompletionCode**

**ConnectionHandle property:**

Read-only. The connection handle for the IBM MQ queue manager object.

**Defined in:**

MQQueueManager class

**Data Type:**

Long

**Syntax:** To get: *hconn*& = *MQQueueManager* .**ConnectionHandle**

**ConnectionStatus property:**

Read-only. Indicates if the object is connected to its queue manager or not.

**Defined in:** MQQueueManager class

**Data Type:** Boolean

**Values:**

- TRUE (-1)
- FALSE (0)

**Syntax:** To get: *status* = *MQQueueManager* .**ConnectionStatus**

**ConnectOptions property:**

Read-Write. These options apply to the Connect method. Initially MQCNO\_NONE.

**Defined in:**

MQQueueManager class

**Data Type:**

Long

**Values:**

- MQCNO\_STANDARD\_BINDING
- MQCNO\_FASTPATH\_BINDING
- MQCNO\_NONE

**Syntax:** To get: *connectoptions*&=MQQueueManager. **ConnectOptions**

To set: MQQueueManager **.ConnectOptions** = *connectoptions*&

**DeadLetterQueueName property:**

Read-only. The MQI DeadLetterQName attribute.

**Defined in:** MQQueueManager class

**Data Type:** String of 48 characters

**Syntax:** to get: *dlqname*\$ = MQQueueManager **.DeadLetterQueueName**

**DefaultTransmissionQueueName property:**

Read-only. The MQI DefXmitQName attribute.

**Defined in:** MQQueueManager class

**Data Type:** String of 48 characters

**Syntax:** To get: *defxmitqname*\$ = MQQueueManager **.DefaultTransmissionQueueName**

**Description property:**

Read-only. The MQI QMgrDesc attribute.

**Defined in:** MQQueueManager class

**Data Type:** String of 64 characters

**Syntax:** To get: *description*\$ = MQQueueManager **.Description**

**DistributionLists property:**

Read-Only. This is the capability of the queue manager to support distribution lists.

**Defined in:**

MQQueueManager class

**Data Type:**

Boolean

**Values:**

- TRUE (-1)
- FALSE (0)

**Syntax:** To get: *distributionlists*= *MQQueueManager*. **DistributionLists**

**InhibitEvent property:**

Read-only. The MQI InhibitEvent attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *inhibevent*& = *MQQueueManager* **.InhibitEvent**

**IsConnected property:**

A value that indicates whether the queue manager is currently connected.

Read-only.

**Defined in:** MQQueueManager class

**Data Type:** Boolean

**Values:**

- TRUE (-1)
- FALSE (0)

**Syntax:** To get: *isconnected* = *MQQueueManager* **.IsConnected**

**IsOpen property:**

A value that indicates whether the queue manager is currently open for inquire.

Read-only.

**Defined in:**

MQQueueManager class

**Data Type:**

Boolean

**Values:**

- TRUE (-1)
- FALSE (0)

**Syntax:** To get: *IsOpen* = *MQQueueManager*. **IsOpen**

**LocalEvent property:**

Read-only. The MQI LocalEvent attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *localevent&* = *MQQueueManager* .**LocalEvent**

**MaximumHandles property:**

Read-only. The MQI MaxHandles attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Syntax:** To get: *maxhandles&* = *MQQueueManager* .**MaximumHandles**

**MaximumMessageLength property:**

Read-only. The MQI MaxMsgLength Queue Manager attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Syntax:** To get: *maxmessagelength&* = *MQQueueManager* .**MaximumMessageLength**

**MaximumPriority property:**

Read-only. The MQI MaxPriority attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Syntax:** To get: *maxpriority& = MQQueueManager .MaximumPriority*

**MaximumUncommittedMessages property:**

Read-only. The MQI MaxUncommittedMsgs attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Syntax:** To get: *maxuncommitted& = MQQueueManager .MaximumUncommittedMessages*

**Name property:**

Read-write. The MQI QMgrName attribute. This property cannot be written once the MQQueueManager is connected.

**Defined in:** MQQueueManager class

**Data Type:** String of 48 characters

**Syntax:** To get: *name\$ = MQQueueManager .name*

To set: *MQQueueManager .name = name\$*

**Note:** Visual Basic reserves the "Name" property for use in the visual interface. Therefore, when using within Visual Basic use lower-case, that is, "name".

**ObjectHandle property:**

Read-only. The object handle for the IBM MQ queue manager object.

**Defined in:**  
MQQueueManager class

**Data type**  
Long

**Syntax:** To get: *hobj& = MQQueueManager .ObjectHandle*

**PerformanceEvent property:**

Read-only. The MQI PerformanceEvent attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *performevent&* = MQQueueManager.PerformanceEvent

**Platform property:**

Read-only. The MQI Platform attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- MQPL\_WINDOWS\_NT
- MQPL\_WINDOWS

**Syntax:** To get: *platform&* = MQQueueManager .Platform

**ReasonCode property:**

Read-only. Returns the reason code set by the last method or property access issued against the object.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- See API reason codes.

**Syntax:** To get: *reasoncode&* = MQQueueManager .ReasonCode



**ReasonName property:**

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC\_QMGR\_NOT\_AVAILABLE".

**Defined in:** MQQueueManager class

**Data Type:** String

**Values:**

- See API reason codes.

**Syntax:** To get: *reasonname\$* = *MQQueueManager* .**ReasonName**

**RemoteEvent property:**

Read-only. The MQI RemoteEvent attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *remoteevent&* = *MQQueueManager* .**RemoteEvent**

**StartStopEvent property:**

Read-only. The MQI StartStopEvent attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *strstpevent&* = *MQQueueManager* .**StartStopEvent**

**SyncPointAvailability property:**

Read-only. The MQI SyncPoint attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Values:**

- MQSP\_AVAILABLE
- MQSP\_NOT\_AVAILABLE

**Syntax:** To get: *syncpointavailability*& = MQQueueManager .SyncPointAvailability

**TriggerInterval property:**

Read-only. The MQI TriggerInterval attribute.

**Defined in:** MQQueueManager class

**Data Type:** Long

**Syntax:** To get: *trigint*& = MQQueueManager .TriggerInterval

**AccessQueue method:**

Creates an MQQueue object and associates it with this MQQueueManager object by setting the connection reference property of the queue. It sets the Name, OpenOptions, DynamicQueueName, and AlternateUserId properties of the MQQueue object to the values provided, and attempts to open it.

If the open is unsuccessful the call fails. An error event is raised against the object. The ReasonCode and CompletionCode, and the MQSession ReasonCode and CompletionCode of the object are set.

The DynamicQueueName, QueueManagerName, and AlternateUserId parameters are optional and default to "".

The OpenOption MQOO\_INQUIRE should be specified in addition to other options if queue properties are to be read.

Do not set the QueueManagerName, or set it to "" if the queue to be opened is local. Otherwise, set it to the name of the remote queue manager that owns the queue, and an attempt is made to open a local definition of the remote queue. For more information about remote queue name resolution and queue manager aliasing, see What are aliases?.

If the Name property is set to a model queue name, specify the name of the dynamic queue to be created in the DynamicQueueName\$ parameter. If the value provided in the DynamicQueueName\$ parameter is "", the value set into the queue object and used on the open call is "AMQ.\*". See "Creating dynamic queues" on page 109 for more information about naming dynamic queues.

**Definition**

**Defined in:** MQQueueManager class.

## Syntax

**Syntax:** set queue = MQQueueManager. **AccessQueue** (Name\$, OpenOptions&, QueueManagerName\$, DynamicQueueName\$, AlternateUserId\$)

## Parameters

*Name\$* String. Name of IBM MQ queue.

*OpenOptions:* Long. Options to be used when queue is opened. See OpenOptions (MQLONG).

*QueueManagerName\$* String. Name of the queue manager that owns the queue to be opened. A value of "" implies the queue manager is local.

*DynamicQueueName\$* String. The name assigned to the dynamic queue at the time the queue is opened when the *Name\$* parameter specifies a model queue.

*AlternateUserId\$* String. The alternate user ID used to validate access when opening the queue.

## AddDistributionList method:

Creates a new MQDistributionList object and sets its connection reference to the owning queue manager.

### Defined in:

MQQueueManager class

**Syntax:** set distributionlist = MQQueueManager. *AddDistributionList*

## Backout method:

Backs out any uncommitted message puts and gets that have occurred as part of a unit of work since the last syncpoint.

**Defined in:** MQQueueManager class

### Syntax:

Call *MQQueueManager.Backout()*

## Begin method:

Begins a unit of work that is coordinated by the queue manager. The begin options affect the behavior of this method.

### Defined in:

MQQueueManager class

### Syntax:

Call *MQQueueManager.Begin()*

**ClearErrorCodes method:**

Resets the CompletionCode to MQCC\_OK and the ReasonCode to MQRC\_NONE for both the MQQueueManager class and the MQSession class.

**Defined in:**

MQQueueManager class

**Syntax:**

Call *MQQueueManager* .**ClearErrorCodes()**

**Commit method:**

Commits any message puts and gets that have occurred as part of a unit of work since the last syncpoint.

**Defined in:** MQQueueManager class

**Syntax:**

Call *MQQueueManager* .Commit()

**Connect method:**

Connects the MQQueueManager object to a real queue manager via the IBM MQ MQI client or server. As well as making the connection, this method also opens the queue manager object so that it can be queried.

Sets IsConnected to TRUE.

A maximum of one MQQueueManager object per ActiveX instance is allowed to connect to a queue manager.

**Defined in:** MQQueueManager class

**Syntax:**

Call *MQQueueManager* .Connect()

**Disconnect method:**

Disconnects the MQQueueManager object from the queue manager.

Sets IsConnected to FALSE.

All Queue objects associated with the MQQueueManager object are made unusable and cannot be re-opened.

Any uncommitted changes (message puts and gets) are committed.

**Defined in:** MQQueueManager class

**Syntax:**

Call *MQQueueManager* .Disconnect()

## MQQueue class

This class represents access to an IBM MQ queue. This connection is provided by an associated MQQueueManager object. When an object of this class is destroyed it is automatically closed.

### Containment

The MQQueue class is contained by the MQQueueManager class.

### Creation

New creates a new MQQueue object and sets all the properties to initial values. Alternatively, use the AccessQueue method of the MQQueueManager class.

### Syntax

```
Dim que As New MQQueue Set que = New MQQueue
```

### Properties

- “AlternateUserId property” on page 711.
- “BackoutRequeueName property” on page 712.
- “BackoutThreshold property” on page 712.
- “BaseQueueName property” on page 712.
- “CloseOptions property” on page 712.
- “CompletionCode property” on page 713.
- “ConnectionReference property” on page 713.
- “CreationDateTime property” on page 713.
- “CurrentDepth property” on page 714.
- “DefaultInputOpenOption property” on page 714.
- “DefaultPersistence property” on page 714.
- “DefaultPriority property” on page 714.
- “DefinitionType property” on page 715.
- “DepthHighEvent property” on page 715.
- “DepthHighLimit property” on page 715.
- “DepthLowEvent property” on page 715.
- “DepthLowLimit property” on page 716.
- “DepthMaximumEvent property” on page 716.
- “DepthHighEvent property” on page 715.
- “DepthHighLimit property” on page 715.
- “DepthLowEvent property” on page 715.
- “DepthLowLimit property” on page 716.
- “DepthMaximumEvent property” on page 716.
- “Description property” on page 716.
- “DynamicQueueName property” on page 716.
- “HardenGetBackout property” on page 717.
- “InhibitGet property” on page 717.
- “InhibitPut property” on page 717.
- “InitiationQueueName property” on page 718.
- “IsOpen property” on page 718.

- “MaximumDepth property” on page 718.
- “MaximumMessageLength property” on page 718.
- “MessageDeliverySequence property” on page 719.
- “ObjectHandle property” on page 719.
- “OpenInputCount property” on page 719.
- “OpenOptions property” on page 720.
- “OpenOutputCount property” on page 720.
- “OpenStatus property” on page 720.
- “ProcessName property” on page 721.
- “QueueManagerName property” on page 721.
- “QueueType Property” on page 721.
- “ReasonCode property” on page 721.
- “ReasonName property” on page 722.
- “RemoteQueueManagerName property” on page 722.
- “RemoteQueueName property” on page 722.
- “ResolvedQueueManagerName property” on page 722.
- “ResolvedQueueName property” on page 723.
- “RetentionInterval property” on page 723.
- “Scope property” on page 723.
- “ServiceInterval property” on page 723.
- “ServiceIntervalEvent property” on page 724.
- “Shareability property” on page 724.
- “TransmissionQueueName property” on page 724.
- “TriggerControl property” on page 724.
- “TriggerData property” on page 725.
- “TriggerDepth property” on page 725.
- “TriggerMessagePriority property” on page 725.
- “TriggerType property” on page 725.
- “Usage property” on page 726.

## Methods

- “ClearErrorCodes method” on page 726
- “Close method” on page 726
- “Get method” on page 726
- “Open method” on page 727
- “Put method” on page 728

## Property Access

If the queue object is not connected to a queue manager, you can read the following properties:

- “CompletionCode property” on page 713
- “OpenStatus property” on page 720
- “ReasonCode property” on page 721

and you can read and write to:

- “AlternateUserId property” on page 711
- “CloseOptions property” on page 712

- “ConnectionReference property” on page 713
- “Name property” on page 719
- “OpenOptions property” on page 720

If the queue object is connected to a queue manager, you can read all the properties.

### Queue Attribute properties

Properties not listed in the previous section are all attributes of the underlying IBM MQ queue. They can be accessed only if the object is connected to a queue manager, and the user's user ID is authorized for Inquire or Set against that queue. If an alternate user ID is set and the current user ID is authorized to use it, the alternate user ID is checked for authorization instead.

The property must be an appropriate property for the given QueueType. See Attributes for queues for more information.

If these conditions do not apply, the property access will set a CompletionCode of MQCC\_FAILED and one of the following ReasonCodes:

- MQRC\_CONNECTION\_BROKEN
- MQRC\_NOT\_AUTHORIZED
- MQRC\_Q\_MGR\_NAME\_ERROR
- MQRC\_Q\_MGR\_NOT\_CONNECTED
- MQRC\_SELECTOR\_NOT\_FOR\_TYPE (CompletionCode is MQCC\_WARNING)

### Opening a queue

The only way to create an MQQueue object is by using the MQQueueManager AccessQueue method or by New. An open MQQueue object remains open (OpenStatus=TRUE) until it is closed or deleted or until the creating queue manager object is deleted or connection is lost to the queue manager. The value of the MQQueue CloseOptions property controls the behavior of the close operation that takes place when the MQQueue object is deleted.

The MQQueueManager AccessQueue method opens the queue using the OpenOptions parameter. The MQQueue.Open method opens the queue using the OpenOptions property. IBM MQ validates the OpenOptions against the user authorization as part of the open queue process.

#### AlternateUserId property:

Read-write. The alternate user ID used to validate access to the queue when it is opened.

This property cannot be set while the object is open (that is, when IsOpen is TRUE).

**Defined in:** MQQueue class

**Data Type:** String of 12 characters

**Syntax:** To get: *altuser\$* = *MQQueue* .**AlternateUserId**

To set: *MQQueue* .**AlternateUserId** = *altuser\$*

**BackoutRequeueName property:**

Read-only. The MQI BackOutRequeueQName attribute.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *backoutrequeuename\$* = *MQQueue* .**BackoutRequeueName**

**BackoutThreshold property:**

Read-only. The MQI BackoutThreshold attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- See BackoutThreshold (MQLONG).

**Syntax:** To get: *backoutthreshold&* = *MQQueue* .**BackoutThreshold**

**BaseQueueName property:**

Read-only. The queue name to which the alias resolves.

Valid only for alias queues.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *baseqname\$* = *MQQueue* .**BaseQueueName**

**CloseOptions property:**

Read-Write. Options used to control what happens when the queue is closed.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQCO\_NONE
- MQCO\_DELETE
- MQCO\_DELETE\_PURGE

MQCO\_DELETE and MQCO\_DELETE\_PURGE are valid only for dynamic queues.

**Syntax:** To get: *closeopt&* = *MQQueue* .**CloseOptions**

To set: *MQQueue* .**CloseOptions** = *closeopt&*



**CompletionCode property:**

Read-only. Returns the completion code set by the last method or property access issued against the object.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQCC\_OK
- MQCC\_WARNING
- MQCC\_FAILED

**Syntax:** To get: *completioncode* = MQQueue .**CompletionCode**

**ConnectionReference property:**

Read-write. Defines the queue manager object to which a queue object belongs. The connection reference cannot be written while a queue is open.

**Defined in:** MQQueue class

**Data Type:** MQQueueManager

**Values:**

- A reference to an active IBM MQ Queue Manager object

**Syntax:** To set: *set MQQueue .ConnectionReference = ConnectionReference*

To get: *set ConnectionReference = MQQueue .ConnectionReference*

**CreationDateTime property:**

Read-only. Date and time this queue was created.

**Defined in:** MQQueue class

**Data Type:** Variant of type 7 (date/time) or EMPTY

**Syntax:** To get: *datetime* = MQQueue .**CreationDateTime**

**CurrentDepth property:**

Read-only. The number of messages currently on the queue.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *currentdepth*& = *MQQueue* .**CurrentDepth**

**DefaultInputOpenOption property:**

Read-only. Controls the way that the queue is opened if the OpenOptions specify MQOO\_INPUT\_AS\_Q\_DEF.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQOO\_INPUT\_EXCLUSIVE
- MQOO\_INPUT\_SHARED

**Syntax:** To get: *defaultinop*& = *MQQueue* .**DefaultInputOpenOption**

**DefaultPersistence property:**

Read-only. The default persistence for messages on a queue.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *defpersistence*& = *MQQueue* .**DefaultPersistence**

**DefaultPriority property:**

Read-only. The default priority for messages on a queue.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *defpriority*& = *MQQueue* .**DefaultPriority**

**DefinitionType property:**

Read-only. Queue definition type.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQQDT\_PREDEFINED
- MQQDT\_PERMANENT\_DYNAMIC
- MQQDT\_TEMPORARY\_DYNAMIC

**Syntax:** To get: *deftype*& = *MQQueue* .**DefinitionType**

**DepthHighEvent property:**

Read-only. The MQI QDepthHighEvent attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *depthhighevent*& = *MQQueue* .**DepthHighEvent**

**DepthHighLimit property:**

Read-only. The MQI QDepthHighLimit attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *depthhighlimit*& = *MQQueue* .**DepthHighLimit**

**DepthLowEvent property:**

Read-only. The MQI QDepthLowEvent attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *depthlowevent*& = *MQQueue* .**DepthLowEvent**

**DepthLowLimit property:**

Read-only. The MQI QDepthLowLimit attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *depthlowlimit*& = MQQueue. **DepthLowLimit**

**DepthMaximumEvent property:**

Read-only. The MQI QDepthMaxEvent attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQEVR\_DISABLED
- MQEVR\_ENABLED

**Syntax:** To get: *depthmaximumevent*& = MQQueue. **DepthMaximumEvent**

**Description property:**

Read-only. A description of the queue.

**Defined in:** MQQueue class

**Data Type:** String of 64 characters

**Syntax:** To get: *description*\$ = MQQueue **.Description**

**DynamicQueueName property:**

Read-write, read-only when the queue is open.

This controls the dynamic queue name used when a model queue is opened. It can be set with a wildcard by the user either as a property set (only when the queue is closed) or as a parameter to MQQueueManager.AccessQueue().

The actual name of the dynamic queue is found by querying QueueName.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Values:**

- Any valid IBM MQ queue name.

**Syntax:** To set: MQQueue **.DynamicQueueName** = *dynamicqueuename*\$

To get: *dynamicqueuename*\$ = MQQueue **.DynamicQueueName**

**HardenGetBackout property:**

Read-only. Whether to maintain an accurate back-out count.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQQA\_BACKOUT\_HARDENED
- MQQA\_BACKOUT\_NOT HARDENED

**Syntax:** To get: *hardengetback&* = *MQQueue* .**HardenGetBackout**

**InhibitGet property:**

Read-write. The MQI InhibitGet attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQQA\_GET\_INHIBITED
- MQQA\_GET\_ALLOWED

**Syntax:** To get: *getstatus&* = *MQQueue* .**InhibitGet**

To set: *MQQueue* .**InhibitGet** = *getstatus&*

**InhibitPut property:**

Read-write. The MQI InhibitPut attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQQA\_PUT\_INHIBITED
- MQQA\_PUT\_ALLOWED

**Syntax:** To get: *putstatus&* = *MQQueue* .**InhibitPut**

To set: *MQQueue* .**InhibitPut** = *putstatus&*

**InitiationQueueName property:**

Read-only. Name of initiation queue.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *initqname\$* = *MQQueue* .**InitiationQueueName**

**IsOpen property:**

Returns whether the queue is open.

Read-only.

**Defined in:** MQQueue class

**Data Type:** Boolean

**Values:**

- TRUE (-1)
- FALSE (0)

**Syntax:** To get: *open* = *MQQueue* .**IsOpen**

**MaximumDepth property:**

Read-only. Maximum queue depth.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *maxdepth&* = *MQQueue* .**MaximumDepth**

**MaximumMessageLength property:**

Read-only. Maximum permitted message length in bytes for this queue.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *maxmlength&* = *MQQueue* .**MaximumMessageLength**

**MessageDeliverySequence property:**

Read-only. Message delivery sequence.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQMDS\_PRIORITY
- MQMDS\_FIFO

**Syntax:** To get: *messdelseq*& = *MQQueue* .**MessageDeliverySequence**

**Name property:**

Read-write. The MQI Queue attribute. This property cannot be written after the MQQueue is open.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *name*\$ = *MQQueue* .**name**

To set: *MQQueue* .**name** = *name*\$

**Note:** Visual Basic reserves the "Name" property for use in the visual interface. Therefore, when using within Visual Basic use lower-case, that is "name".

**ObjectHandle property:**

Read-only. The object handle for the IBM MQ queue object.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *hobj*& = *MQQueue* .**ObjectHandle**

**OpenInputCount property:**

Read-only. Number of opens for input.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get:

*openincount*& = *MQQueue* .**OpenInputCount**

**OpenOptions property:**

Read-write. Options to be used for opening the queue.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- See OpenOptions (MQLONG).

**Syntax:** To get:

`openopt& = MQQueue.OpenOptions`

To set: `MQQueue.OpenOptions = openopt&`

**OpenOutputCount property:**

Read-only. Number of opens for output.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get:

`openoutputcount& = MQQueue.OpenOutputCount`

**OpenStatus property:**

Read-only. Indicates if the queue is opened or not. Initial value is TRUE after AccessQueue method or FALSE after New.

**Defined in:** MQQueue class

**Data Type:** Boolean

**Values:**

- TRUE (-1)
- FALSE (0)

**Syntax:** To get:

`status& = MQQueue.OpenStatus`



**ProcessName property:**

Read-only. The MQI ProcessName attribute.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *procname\$* = *MQQueue* .**ProcessName**

**QueueManagerName property:**

Read-write. The IBM MQ queue manager name.

**Defined in:** MQQueue class

**Data Type:** String

**Syntax:** To get: *QueueManagerName\$* = *MQQueue* .**QueueManagerName**

To set: *MQQueue* .**QueueManagerName** = *QueueManagerName\$*

**QueueType Property:**

Read-only. The MQI QType attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQQT\_ALIAS
- MQQT\_LOCAL
- MQQT\_MODEL
- MQQT\_REMOTE

**Syntax:** To get: *queuetype&* = *MQQueue* .**QueueType**

**ReasonCode property:**

Read-only. Returns the reason code set by the last method or property access issued against the object.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- See API reason codes.

**Syntax:** To get: *reasoncode&* = *MQQueue* .**ReasonCode**

**ReasonName property:**

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC\_QMGR\_NOT\_AVAILABLE".

**Defined in:** MQQueue class

**Data Type:** String

**Values:**

- See API reason codes.

**Syntax:** To get: *reasonname\$* = *MQQueue* .ReasonName

**RemoteQueueManagerName property:**

Read-only. Name of remote queue manager. Valid for remote queues only.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *remqmanname\$* = *MQQueue* .RemoteQueueManagerName

**RemoteQueueName property:**

Read-only. The name of the queue as it is known on the remote queue manager. Valid for remote queues only.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *remqname\$* = *MQQueue* .RemoteQueueName

**ResolvedQueueManagerName property:**

Read-only. The name of the final destination queue manager as known to the local queue manager.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *resqmanname\$* = *MQQueue* .ResolvedQueueManagerName

**ResolvedQueueName property:**

Read-only. The name of the final destination queue as known to the local queue manager.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *resqname\$* = *MQQueue* .**ResolvedQueueName**

**RetentionInterval property:**

Read-only. The period of time for which the queue should be retained.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *retinterval&* = *MQQueue* .**RetentionInterval**

**Scope property:**

Read-only. Controls whether an entry for this queue also exists in a cell directory.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQSCO\_Q\_MGR
- MQSCO\_CELL

**Syntax:** To get: *scope&* = *MQQueue* .**Scope**

**ServiceInterval property:**

Read-only. The MQI QServiceInterval attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Syntax:** To get: *serviceinterval&* = *MQQueue* .**ServiceInterval**

**ServiceIntervalEvent property:**

Read-only. The MQI QServiceIntervalEvent attribute.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQQSIE\_HIGH
- MQQSIE\_OK
- MQQSIE\_NONE

**Syntax:** To get: *serviceintvalevent& = MQQueue.ServiceIntervalEvent*

**Shareability property:**

Read-only. Queue shareability.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQQA\_SHAREABLE
- MQQA\_NOT\_SHAREABLE

**Syntax:** To get: *shareability& = MQQueue.Shareability*

**TransmissionQueueName property:**

Read-only. Transmission queue name. Valid for remote queues only.

**Defined in:** MQQueue class

**Data Type:** String of 48 characters

**Syntax:** To get: *transqname\$ = MQQueue.TransmissionQueueName*

**TriggerControl property:**

Read-write. Trigger control.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQTC\_OFF
- MQTC\_ON

**Syntax:** To get: *trigcontrol& = MQQueue.TriggerControl*

To set: *MQQueue.TriggerControl = trigcontrol&*

**TriggerData property:**

Read-write. Trigger data.

**Defined in:** MQTTQueue class

**Data Type:** String of 64 characters

**Syntax:** To get: *trigdata\$* = MQTTQueue .TriggerData

To set: MQTTQueue .TriggerData = *trigdata\$*

**TriggerDepth property:**

Read-write. The number of messages that have to be on the queue before a trigger message is written.

**Defined in:** MQTTQueue class

**Data Type:** Long

**Syntax:** To get: *trigdepth&* = MQTTQueue .TriggerDepth

To set: MQTTQueue .TriggerDepth = *trigdepth&*

**TriggerMessagePriority property:**

Read-write. Threshold message priority for triggers.

**Defined in:** MQTTQueue class

**Data Type:** Long

**Syntax:** To get: *trigmesspriority&* = MQTTQueue .TriggerMessagePriority

To set: MQTTQueue .TriggerMessagePriority = *trigmesspriority&*

**TriggerType property:**

Read-write. Trigger type.

**Defined in:** MQTTQueue class

**Data Type:** Long

**Values:**

- MQTT\_NONE
- MQTT\_FIRST
- MQTT EVERY
- MQTT\_DEPTH

**Syntax:** To get: *trigtype&* = MQTTQueue .TriggerType

To set: MQTTQueue .TriggerType = *Trigtype&*

**Usage property:**

Read-only. Indicates what the queue is used for.

**Defined in:** MQQueue class

**Data Type:** Long

**Values:**

- MQUS\_NORMAL
- MQUS\_TRANSMISSION

**Syntax:** To get: *usage* = *MQQueue* .Usage

**ClearErrorCodes method:**

Resets the CompletionCode to MQCC\_OK and the ReasonCode to MQRC\_NONE for both the MQQueue class and the MQSession class.

**Defined in:** MQQueue class

**Syntax:**

Call *MQQueue* .ClearErrorCodes()

**Close method:**

Closes a queue using the current values of CloseOptions.

**Defined in:** MQQueue class

**Syntax:**

Call *MQQueue* .Close()

**Get method:**

Retrieves a message from the queue.

This method takes an MQMessage object as a parameter, using some of the fields in the MQMD of the object as input parameters. In particular, the MessageId and CorrelId fields are used, so it is important to ensure that these fields are set as required. For more information about these fields, see MsgId (MQBYTE24) and CorrelId (MQBYTE24).

If the method fails then the MQMessage object is unchanged. If it succeeds, the MQMD and Message Data portions of the MQMessage object are replaced by the MQMD and Message Data from the incoming message. The MQMessage control properties are set as follows

- **MessageLength** is set to length of the IBM MQ message
- **DataLength** is set to length of the IBM MQ message
- **DataOffset** is set to zero

**Defined in:**

MQQueue class

**Syntax:**

Call *MQQueue* .Get(*Message*, *GetMsgOptions*, *GetMsgLength*)

## Parameters

Message:

MQMessage Object representing message to be retrieved.

GetMsgOptions:

Optional MQGetMessageOptions object to control the get operation. If this parameter is not specified, default MQGetMessageOptions are used.

GetMsgLength:

Optional 2 or 4 byte length value to control the maximum length of the IBM MQ message that is retrieved from the queue.

If the MQGMO\_ACCEPT\_TRUNCATED\_MSG option is specified, the GET succeeds with a completion code of MQCC\_WARNING and a reason code of MQRC\_TRUNCATED\_MSG\_ACCEPTED if the message size exceeds the specified length.

The MessageData holds the first GetMsgLength bytes of data.

If MQGMO\_ACCEPT\_TRUNCATED\_MSG *is not* specified, and the message size exceeds the specified length, the completion code of MQCC\_FAILED together with reason code MQRC\_TRUNCATED\_MESSAGE\_FAILED is returned.

If the contents of the message buffer are undefined, the total message length is set to the full length of the message that would have been retrieved.

If the message length parameter is not specified, the length of the message buffer is automatically adjusted to at least the size of the incoming message.

## Open method:

Opens a queue using the current values of:

1. QueueName
2. QueueManagerName
3. AlternateUserId
4. DynamicQueueName

## Defined in:

MQQueue class

## Syntax:

Call `MQQueue.Open()`

**Put method:**

Places a message onto the queue.

This method takes an `MQMessage` object as a parameter. The Message Descriptor (MQMD) properties of this object might be altered as a result of this method. The values they have immediately after this method has run are the values that were put onto the IBM MQ.

Modifications to the `MQMessage` object after the Put has completed do not affect the actual message on the IBM MQ queue.

**Defined in:**

`MQQueue` class

**Syntax:**

Call `MQQueue.Put(Message, PutMsgOptions)`

**Parameters**

Message

`MQMessage` object representing message to be put.

`PutMsgOptions`

`MQPutMessageOptions` object containing options to control the put operation. If these are not specified, default `PutMessageOptions` are used.

**MQMessage class**

This class represents an IBM MQ message. It includes properties to encapsulate the IBM MQ message descriptor (MQMD), and provides a buffer to hold the application-defined message data.

The class includes Write methods to copy data from an ActiveX application to an `MQMessage` object. Similarly, the class includes Read methods to copy data from an `MQMessage` object to an ActiveX application. The class manages the allocation and deallocation of memory for the buffer automatically. The application does not have to declare the size of the buffer when an `MQMessage` object is created because the buffer grows to accommodate data written to it.

You cannot place a message onto an IBM MQ queue if the buffer size exceeds the `MaximumMessageLength` property of that queue.

After it has been constructed, an `MQMessage` object can be Put onto an IBM MQ queue using the `MQQueue.Put` method. This method takes a copy of the MQMD and message data portions of the object and places that copy on the queue. The application can therefore modify or delete an `MQMessage` object after the Put, without affecting the message on the IBM MQ queue. The queue manager can adjust some of the fields in the MQMD when it copies the message on the IBM MQ queue.

An incoming message can be read into an `MQMessage` object using the `MQQueue.Get` method. This replaces any MQMD or message data that might already have been in the `MQMessage` object with values from the incoming message. It adjusts the size of the data buffer of the `MQMessage` object to match the size of the incoming message data.

**Containment**

Messages are contained by the `MQSession` class.



## Creation

**New** creates an MQMessage object. Its Message Descriptor properties are initially set to default values, and its Message Data buffer is empty.

## Syntax

```
Dim msg As New MQMessage
```

or

```
Set msg = New MQMessage
```

## Properties

The control properties are:

- “CompletionCode property” on page 731
- “DataLength property” on page 732
- “DataOffset property” on page 732
- “MessageLength property” on page 733
- “ReasonCode property” on page 733
- “ReasonName property” on page 733

The Message Descriptor properties are:

- “AccountingToken property” on page 734
- “AccountingTokenHex property” on page 734
- “ApplicationIdData property” on page 734
- “ApplicationOriginData property” on page 735
- “BackoutCount property” on page 735
- “CharacterSet property” on page 735
- “CorrelationId property” on page 736
- “CorrelationIdHex property” on page 736
- “Encoding property” on page 737
- “Expiry property” on page 738
- “Feedback property” on page 738
- “Format property” on page 738
- “GroupId property” on page 739
- “GroupIdHex property” on page 739
- “MessageData property” on page 740
- “MessageFlags property” on page 740
- “MessageId property” on page 740
- “MessageIdHex property” on page 741
- “MessageSequenceNumber property” on page 741
- “MessageType property” on page 742
- “Offset property” on page 742
- “OriginalLength property” on page 742
- “Persistence property” on page 743
- “Priority property” on page 743
- “PutApplicationName property” on page 743
- “PutApplicationType property” on page 744

- “PutDateTime property” on page 744
- “ReplyToQueueManagerName property” on page 744
- “ReplyToQueueName property” on page 745
- “Report property” on page 745
- “TotalMessageLength property” on page 745
- “UserId property” on page 746

## Methods

- “ClearErrorCodes method” on page 746
- “ClearMessage method” on page 746
- “Read method” on page 746
- “ReadBoolean method” on page 747
- “ReadByte method” on page 747
- “ReadDecimal2 method” on page 747
- “ReadDecimal4 method” on page 747
- “ReadDouble method” on page 748
- “ReadDouble4 method” on page 748
- “ReadFloat method” on page 749
- “ReadInt2 method” on page 749
- “ReadInt4 method” on page 749
- “ReadLong method” on page 749
- “ReadNullTerminatedString method” on page 750
- “ReadShort method” on page 750
- “ReadString method” on page 750
- “ReadUInt2 method” on page 751
- “ReadUnsignedByte method” on page 751
- “ReadUTF method” on page 751
- “ResizeBuffer method” on page 752
- “Write method” on page 752
- “WriteBoolean method” on page 753
- “WriteByte method” on page 753
- “WriteDecimal2 method” on page 753
- “WriteDecimal4 method” on page 754
- “WriteDouble method” on page 754
- “WriteDouble4 method” on page 754
- “WriteFloat method” on page 755
- “WriteInt2 method” on page 755
- “WriteInt4 method” on page 755
- “WriteLong method” on page 755
- “WriteNullTerminatedString method” on page 756
- “WriteShort method” on page 756
- “WriteString method” on page 756
- “WriteUInt2 method” on page 757
- “WriteUnsignedByte method” on page 757
- “WriteUTF method” on page 757

## Property access

All properties can be read at any time.

The control properties are read-only, except for `DataOffset` which is read-write. The Message Descriptor properties are all read-write, except `BackoutCount` and `TotalMessageLength` which are both read-only.

Note however that some of the MQMD properties might be modified by the queue manager when the message is put onto an IBM MQ queue. See the fields in MQMD for details of how they might be modified.

## Data conversion

You can pass binary data to an IBM MQ message by setting the **CharacterSet** property to match the Coded Character Set Identifier of the queue manager (MQCCSI\_Q\_MGR), and passing the data to the message as a string. If the string needs to include Unicode or ASCII code numbers, you can use the `chr$` function to convert them to string format.

The Read and Write methods perform data conversion. They convert between the ActiveX internal formats, and the IBM MQ message formats as defined by the Encoding and CharacterSet properties from the message descriptor. When writing a message, set values into Encoding and CharacterSet that match the characteristics of the recipient of the message before issuing a Write method. When reading a message, this step is not normally required because these values will have been set from those in the incoming MQMD.

This is an additional data conversion step that happens after any conversion performed by the `MQQueue.Get` method.

### CompletionCode property:

Read-only. Returns the IBM MQ completion code set by the most recent method or property access issued against this object.

**Defined in:** MQMessage class

**Data Type:** Long

#### Values:

- MQCC\_OK
- MQCC\_WARNING
- MQCC\_FAILED

**Syntax:** To get: `completioncode& = MQMessage .CompletionCode`

**DataLength property:**

Read-only. This property returns the value:

`MQMessage.MessageLength - MQMessage.DataOffset`

It can be used before a Read method, to check that the expected number of characters are actually present in the buffer.

The initial value is zero.

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *bytesleft* = *MQMessage* .**DataLength**

**DataOffset property:**

Read-write. The current position within the Message Data portion of the message object.

The value is expressed as a byte offset from the start of the message data buffer; the first character in the buffer corresponds to a DataOffset value of zero.

A read or write method commences its operation at the character referenced by DataOffset. These methods process data in the buffer sequentially from this position, and update DataOffset to point to the byte (if any) immediately following the last byte processed.

DataOffset can take only values in the range zero to MessageLength inclusive. When DataOffset = MessageLength it is pointing to the end, that is the first invalid character of the buffer. Write methods are permitted in this situation - they extend the data in the buffer and increase MessageLength by the number of bytes added. Reading beyond the end of the buffer is not valid.

The initial value is zero.

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *currpos* = *MQMessage* .**DataOffset**

To set: *MQMessage* .**DataOffset** = *currpos*

**MessageLength property:**

Read-only. Returns the total length of the Message Data portion of the message object in characters, irrespective of the value of DataOffset.

The initial value is zero. It is set to the incoming Message Length after a Get method invocation that referenced this message object. It is incremented if the application uses a Write method to add data to the object. It is unaffected by Read methods.

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *msglength&* = *MQMessage* .**MessageLength**

**ReasonCode property:**

Read-only. Returns the reason code set by the most recent method or property access issued against this object.

**Defined in:** MQMessage class

**Data Type:** Long

**Values:**

- See API reason codes.

**Syntax:** To get: *reasoncode&* = *MQMessage* .**ReasonCode**

**ReasonName property:**

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC\_QMGR\_NOT\_AVAILABLE". **Defined in:** MQMessage class

**Data Type:** String

**Values:**

- See API reason codes.

**Syntax:** To get: *reasonname\$* = *MQMessage* .**ReasonName**

**AccountingToken property:**

Read-write. The MQMD AccountingToken - part of the message Identity Context.

Its initial value is all nulls.

**Defined in:** MQMessage class

**Data Type:** String of 32 characters

**Syntax:** To get: *actoken\$* = *MQMessage* .**AccountingToken**

To set: *MQMessage* . **AccountingToken** = *actoken\$*

See "Message Descriptor properties" on page 683 for more information about when you must use AccountingTokenHex in place of the AccountingToken property.

**AccountingTokenHex property:**

Read-write. The MQMD AccountingToken - part of the message Identity Context.

Every two characters represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B", and so on.

You must supply 64 valid hexadecimal characters.

Its initial value is "0...0"

**Defined in:** MQMessage class

**Data Type:** String of 64 hexadecimal characters representing 32 ASCII characters

**Syntax:** To get: *actokenh\$* = *MQMessage* .**AccountingTokenHex**

To set: *MQMessage* . **AccountingTokenHex** = *actokenh\$*

See "Message Descriptor properties" on page 683 for more information about when you must use AccountingTokenHex in place of the AccountingToken property.

**ApplicationIdData property:**

Read-write. The MQMD AppllIdentityData - part of the message Identity Context.

Its initial value is all blanks.

**Defined in:** MQMessage class

**Data Type:** String of 32 characters

**Syntax:** To get: *applid\$* = *MQMessage* .**ApplicationIdData**

To set: *MQMessage* .**ApplicationIdData** = *applid\$*

**ApplicationOriginData property:**

Read-write. The MQMD ApplOriginData - part of the message origin context.

Its initial value is all blanks.

**Defined in:** MQMessage class

**Data Type:** String of 4 characters

**Syntax:** To get: *applor\$* = MQMessage .ApplicationOriginData

To set: MQMessage .ApplicationOriginData = *applor\$*

**BackoutCount property:**

Read-only. The MQMD BackoutCount.

Its initial value is 0

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *backoutct&* = MQMessage .BackoutCount

**CharacterSet property:**

Read-write. The MQMD CodedCharSetId.

Its initial value is the special value **MQCCSI\_Q\_MGR**.

If the CharacterSet is set to **MQCCSI\_Q\_MGR**, the code page for the current locale is used for character conversion in the WriteString method. For server applications, the code page used is the code page of the queue manager. For client applications, it is the default current locale code page.

For example:

```
msg.CharacterSet = MQCCSI_Q_MGR
msg.WriteString(chr$(n))
```

where 'n' is greater than or equal to zero and less than or equal to 255, results in a single byte of value of 'n' being written to the buffer.

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *:30ccid&* = MQMessage .CharacterSet

To set: MQMessage .CharacterSet = *ccid&*

**Example**

If you want the string written out in code page 437, issue:

```
Message.CharacterSet = 437
Message.WriteString ("string to be written")
```

Set the value you want in the CharacterSet before issuing any WriteString calls.

**CorrelationId property:**

Read-write. The CorrelationId to be included in the MQMD of a message when put on a queue. Also the ID to be matched against when getting a message from a queue.

Its initial value is null.

**Defined in:** MQMessage class

**Data Type:** String of 24 characters

**Syntax:** To get: *correlid\$ = MQMessage .CorrelationId* To set: *MQMessage .CorrelationId = correlid\$*

See "Message Descriptor properties" on page 683 for more information about when you must use CorrelationIdHex in place of the CorrelationId property.

**CorrelationIdHex property:**

Read-write. The CorrelationId to be included in the MQMD of a message when put on a queue. Also the CorrelationId to be matched against when getting a message from a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B", and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0...0".

**Defined in:** MQMessage class

**Data Type:** String of 48 hexadecimal characters representing 24 ASCII characters

**Syntax:** To get: *correlidh\$ = MQMessage .CorrelationIdHex*

To set: *MQMessage .CorrelationIdHex = correlidh\$*

See "Message Descriptor properties" on page 683 for a discussion of when you must use CorrelationIdHex in place of the CorrelationId property.



### Encoding property:

Read-write. The MQMD field that identifies the representation used for numeric values in the application message data.

Its initial value is the special value MQENC\_NATIVE, which varies by platform.

This property is used by the following methods:

- ReadDecimal2 method
- ReadDecimal4 method
- ReadDouble method
- ReadDouble4 method
- ReadFloat method
- ReadInt2 method
- ReadInt4 method
- ReadLong method
- ReadShort method
- ReadUInt2 method
- WriteDecimal2 method
- WriteDecimal4 method
- WriteDouble method
- WriteDouble4 method
- WriteFloat method
- WriteInt2 method
- WriteInt4 method
- WriteLong method
- WriteShort method
- WriteUInt2 method

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *encoding&* = MQMessage .**Encoding** To set: MQMessage .**Encoding** = *encoding&*

If you are preparing to write data to the message buffer, you should set this field to match the characteristics of the receiving queue manager platform if the receiving queue manager is incapable of performing its own data conversion.

**Expiry property:**

Read-write. The MQMD expiry time field, expected in tenths of a second.

Its initial value is the special value MQEL\_UNLIMITED

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *expiry&* = MQMessage **.Expiry**

To set: MQMessage **.Expiry** = *expiry&*

**Feedback property:**

Read-write. The MQMD feedback field.

Its initial value is the special value MQFB\_NONE.

**Defined in:** MQMessage class

**Data Type:** Long

**Values:**

- See Feedback.

**Syntax:** To get: *feedback&* = MQMessage **.Feedback**

To set: MQMessage **.Feedback** = *feedback&*

**Format property:**

Read-write. The MQMD format field. Gives the name of a built-in or user-defined format that describes the nature of the Message Data.

Its initial value is the special value MQFMT\_NONE.

**Defined in:** MQMessage class

**Data Type:** String of 8 characters

**Syntax:** To get: *format\$* = MQMessage **.Format**

To set: MQMessage **.Format** = *format\$*

**GroupId property:**

Read-write. The GroupId to be included in the MQPMR of a message when put on a queue. Also the ID to be matched against when getting a message from a queue. Its initial value is all nulls.

**Defined in:**

MQMessage class

**Data Type:**

String of 24 characters

**Syntax:** To get: *groupid\$* = *MQMessage.GroupId*

To set: *MQMessage.GroupId* = *groupid\$*

See "Message Descriptor properties" on page 683 for more information about when you must use GroupIdHex in place of the GroupId property.

**GroupIdHex property:**

Read-write. The GroupId to be included in the MQPMR of a message when put on a queue. Also the ID to be matched against when getting a message from a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0...0".

**Defined in:**

MQMessage class

**Data Type:**

String of 48 hexadecimal characters representing 24 ASCII characters.

**Syntax:** To get: *groupidh\$* = *MQMessage.GroupIdHex*

To set: *MQMessage.GroupIdHex* = *groupidh\$*

See "Message Descriptor properties" on page 683 for more information about when you must use GroupIdHex in place of the GroupId property.

**MessageData property:**

Read-write. Retrieves or sets the entire contents of a message as a character string.

**Defined in:** MQMessage class

**Data Type:** Variant

**Note:** The data type used by this property is Variant but MQAX expects this to be a variant type of String. If you pass in a variant of other than this type then the error MQRC\_OBJECT\_TYPE\_ERROR will be returned.

**Syntax:** To get: *String\$* = MQMessage .**MessageData**

To set: MQMessage .**MessageData** = *String\$*

**MessageFlags property:**

Read-Write. Message flags specifying Segmentation control information. The initial value is 0.

**Defined in:**

MQMessage class

**Data Type:**

Long

**Values:**

See MsgFlags (MQLONG).

**Syntax:** To get: *messageflags&* = MQMessage. **MessageFlags**

To set: MQMessage. **MessageFlags** = *messageflags&*

**MessageId property:**

Read-write. The MessageId to be included in the MQMD of a message when put on a queue. Also the ID to be matched against when getting a message from a queue.

Its initial value is all nulls.

**Defined in:** MQMessage class

**Data Type:** String of 24 characters

**Syntax:** To get: *messageid\$* = MQMessage .**MessageId**

To set: MQMessage .**MessageId** = *messageid\$*

See "Message Descriptor properties" on page 683 for more information about when you must use MessageIdHex in place of the MessageId property.

**MessageIdHex property:**

Read-write. The MessageId to be included in the MQMD of a message when put on a queue. Also the MessageId to be matched against when getting a message from a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B", and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0...0".

**Defined in:** MQMessage class

**Data Type:** String of 48 hexadecimal characters representing 24 ASCII characters

**Syntax:** To get: *messageidh\$* = *MQMessage* .**MessageIdHex**

To set: *MQMessage* .**MessageIdHex** = *messageidh\$*

See "Message Descriptor properties" on page 683 for more information about when you must use MessageIdHex in place of the MessageId property.

**MessageSequenceNumber property:**

Read-Write. Sequence information identifying a message within a group. The initial value is 1.

**Defined in:**  
MQMessage class

**Data Type:**  
Long

**Values:**  
See MsgSeqNumber (MQLONG).

**Syntax:** To get: *sequencenumber&* = *MQMessage* .**SequenceNumber**

To set: *MQMessage* .**SequenceNumber** = *sequencenumber&*

**MessageType property:**

Read-write. The MQMD MsgType field.

Its initial value is MQMT\_DATAGRAM.

**Defined in:** MQMessage class

**Data Type:** Long

**Values:**

- See MsgType (MQLONG).

**Syntax:** To get: *msgtype&* = *MQMessage* .**MessageType**

To set: *MQMessage* .**MessageType** = *msgtype&*

**Offset property:**

Read-Write. The offset in a segmented message. The initial value is 0.

**Defined in:**

MQMessage class

**Data Type:**

Long

**Values:**

See Offset (MQLONG).

**Syntax:** To get: *offset&* = *MQMessage* .**Offset**

To set: *MQMessage* .**Offset** = *offset&*

**OriginalLength property:**

Read-Write. The original length of a segmented message. The initial value is MQOL\_UNDEFINED.

**Defined in:**

MQMessage class

**Data Type:**

Long

**Values:**

See OriginalLength (MQLONG).

**Syntax:** To get: *originallength&* = *MQMessage* .**OriginalLength**

To set: *MQMessage* .**OriginalLength** = *originallength&*

**Persistence property:**

Read-write. The message's persistence setting.

Its initial value is MQPER\_PERSISTENCE\_AS\_Q\_DEF.

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *persist&* = MQMessage **.Persistence**

To set: MQMessage **.Persistence** = *persist&*

**Priority property:**

Read-write. The message's priority.

Its initial value is the special value MQPRI\_PRIORITY\_AS\_Q\_DEF

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *priority&* = MQMessage **.Priority**

To set: MQMessage **.Priority** = *priority&*

**PutApplicationName property:**

Read-write. The MQMD PutAppNm - part of the Message Origin context.

Its initial value is all blanks.

**Defined in:** MQMessage class

**Data Type:** String of 28 characters

**Syntax:** To get: *putapplnm\$* = MQMessage **.PutApplicationName**

To set: MQMessage **.PutApplicationName** = *putapplnm\$*

**PutApplicationType property:**

Read-write. The MQMD PutApplType - part of the Message Origin context.

Its initial value is MQAT\_NO\_CONTEXT

**Defined in:** MQMessage class

**Data Type:** Long

**Values:**

- See PutApplType (MQLONG).

**Syntax:** To get: *putappltp&* = MQMessage .**PutApplicationType**

To set: MQMessage .**PutApplicationType** = *putappltp&*

**PutDateTime property:**

Read/write. This property combines the MQMD PutDate and PutTime fields. These are parts of the Message Origin context that indicate when the message was put.

The ActiveX Extension converts between ActiveX date/time format and the Date and Time formats used in an IBM MQ MQMD. If a message is received which has an invalid PutDate or PutTime, then the PutDateTime property after the get method is set to EMPTY.

Its initial value is EMPTY.

**Defined in:** MQMessage class

**Data Type:** Variant of type 7 (date/time) or EMPTY.

**Syntax:** To get: *datetime* = MQMessage .**PutDateTime**

To set: MQMessage .**PutDateTime** = *datetime*

**ReplyToQueueManagerName property:**

Read-write. The MQMD ReplyToQMgr field.

Its initial value is all blanks

**Defined in:** MQMessage class

**Data Type:** String of 48 characters

**Syntax:** To get: *replytoqmgr\$* = MQMessage .**ReplyToQueueManagerName**

To set: MQMessage .**ReplyToQueueManagerName** = *replytoqmgr\$*



**ReplyToQueueName property:**

Read-write. The MQMD ReplyToQ field.

Its initial value is all blanks

**Defined in:** MQMessage class

**Data Type:** String of 48 characters

**Syntax:** To get: *replytoq\$* = MQMessage **.ReplyToQueueName**

To set: MQMessage **.ReplyToQueueName** = *replytoq\$*

**Report property:**

Read-write. The message's Report options.

Its initial value is MQRO\_NONE.

**Defined in:** MQMessage class

**Data Type:** Long

**Values:**

- See Report.

**Syntax:** To get: *report&* = MQMessage **.Report**

To set: MQMessage **.Report** = *report&*

**TotalMessageLength property:**

Read-only. Retrieves the length of the last message received by MQGET. If the message has not been truncated, this value is equal to the value of the MessageLength property.

**Defined in:** MQMessage class

**Data Type:** Long

**Syntax:** To get: *totalmessagelength&* = MQMessage **.TotalMessageLength**

**UserId property:**

Read-write. The MQMD UserIdentifier - part of the message Identity Context.

Its initial value is all blanks.

**Defined in:** MQMessage class

**Data Type:** String of 12 characters

**Syntax:** To get: *userid\$* = MQMessage .**UserId**

To set: MQMessage .**UserId** = *userid\$*

**ClearErrorCodes method:**

Resets the CompletionCode to MQCC\_OK and the ReasonCode to MQRC\_NONE for both the MQMessage class and the MQSession class.

**Defined in:** MQMessage class

**Syntax:**

Call *MQMessage*.ClearErrorCodes()

**ClearMessage method:**

This method clears the data buffer portion of the MQMessage object. Any Message Data in the data buffer is lost, because MessageLength, DataLength, and DataOffset are all set to zero.

The Message Descriptor (MQMD) portion is unaffected; an application might need to modify some of the MQMD fields before reusing the MQMessage object. To set the MQMD fields back use New to replace the object with a new instance.

**Defined in:** MQMessage class

**Syntax:**

Call *MQMessage*.ClearMessage()

**Read method:**

Reads a sequence of bytes from the message buffer into a byte array. DataOffset is incremented and Data Length decremented by the number of bytes read.

**Defined in:**

MQMessage class

**Syntax:** Data = MQMessage. **Read** (len&)

**Parameters:**

*len&*: Long. Length of data in bytes to be read.

**ReadBoolean method:**

Reads a 1-byte Boolean value from the current position in the message buffer and returns a 2-byte Boolean TRUE(-1)/FALSE(0) value. DataOffset is incremented by one and Data Length is decremented by one.

**Defined in:**

MQMessage class

**Syntax:** *value* = MQMessage.**ReadBoolean**

**ReadByte method:**

Starting with the byte referred to by DataOffset, the ReadByte method reads 1 byte from the Message Data buffer and returns it as an Integer (signed 2-byte) integer value in the range -128 to 127.

The method fails if MQMessage.DataLength is less than 1 when it is issued.

DataOffset is incremented by 1 and DataLength is decremented by 1 if the method succeeds.

The byte of message data is assumed to be a signed binary integer.

**Defined in:**

MQMessage class

**Syntax:**

*integer%* = MQMessage.**ReadByte**

**ReadDecimal2 method:**

Reads a 2-byte packed decimal number and returns it as a signed 2-byte integer value. DataOffset is incremented by two and Data Length is decremented by two.

**Defined in:**

MQMessage class

**Syntax:** *value%* = MQMessage.**ReadDecimal2**

**ReadDecimal4 method:**

Reads a 4-byte packed decimal number and returns it as a signed 4-byte integer value. DataOffset is incremented by four and Data Length is decremented by four.

**Defined in:**

MQMessage class

**Syntax:**

Call *value&* = MQMessage.**ReadDecimal4**

### **ReadDouble method:**

Starting with the byte referred to by `DataOffset`, the `ReadDouble` method reads 8 bytes from the Message Data buffer and returns them as a `Double` (signed 8-byte) floating point value.

The method fails if `MQMessage.DataLength` is less than 8 when it is issued.

`DataOffset` is incremented by 8 and `DataLength` is decremented by 8 if the method succeeds.

The 8 characters of message data are assumed to be a binary floating point number. The encoding is specified by the `MQMessage.Encoding` property. Note that conversion from System/360 format is not supported.

#### **Defined in:**

`MQMessage` class

#### **Syntax:**

*doublev#* = *MQMessage*.**ReadDouble**

### **ReadDouble4 method:**

The `ReadDouble4` and `WriteDouble4` methods are alternatives to `ReadFloat` and `WriteFloat`. This is because they support 4-byte System/390 floating point message values that are too large to convert to 4-byte IEEE floating point format.

Starting with the byte referred to by `DataOffset`, the `ReadDouble4` method reads 4 bytes from the Message Data buffer and returns them as a `Double` (signed 8-byte) floating point value.

The method fails if `MQMessage.DataLength` is less than 4 when it is issued.

`DataOffset` is incremented by 4 and `DataLength` is decremented by 4 if the method succeeds.

The 4 characters of message data are assumed to be a binary floating point number. The encoding is specified by the `MQMessage.Encoding` property. Note that conversion from System/360 format is not supported.

#### **Defined in:**

`MQMessage` class

#### **Syntax:**

*doublev#* = *MQMessage*.**ReadDouble4**

**ReadFloat method:**

Starting with the byte referred to by `DataOffset`, the `ReadFloat` method reads 4 bytes from the Message Data buffer and returns them as a `Single` (signed 4-byte) floating point value.

The method fails if `MQMessage.DataLength` is less than 4 when it is issued.

`DataOffset` is incremented by 4 and `DataLength` is decremented by 4 if the method succeeds.

The 4 characters of message data are assumed to be a floating point number. The encoding is specified by the `MQMessage.Encoding` property. Note that conversion from System/360 format is not supported.

**Defined in:**

`MQMessage` class

**Syntax:**

*single!* = `MQMessage` .**ReadFloat**

**ReadInt2 method:**

The method is identical to the `ReadShort` method.

**Syntax:**

*integer%* = `MQMessage` .**ReadInt2**

**ReadInt4 method:**

This method is identical to the `ReadLong` method.

**Syntax:**

*bigint&* = `MQMessage` .**ReadInt4**

**ReadLong method:**

Starting with the byte referred to by `DataOffset`, the `ReadLong` method reads 4 bytes from the Message Data buffer and returns them as a `Long` (signed 4-byte) integer value.

The method fails if `MQMessage.DataLength` is less than 4 when it is issued.

`DataOffset` is incremented by 4 and `DataLength` is decremented by 4 if the method succeeds.

The 4 characters of message data are assumed to be a binary integer. The encoding is specified by the `MQMessage.Encoding` property.

**Defined in:**

`MQMessage` class

**Syntax:**

*bigint&* = `MQMessage` .**ReadLong**

### **ReadNullTerminatedString method:**

This method is for use in place of ReadString if the string might contain embedded null characters.

This method reads the specified number of bytes from the message data buffer starting with the byte referred to by DataOffset and returns it as an ActiveX string. If the string contains an embedded null before the end then the length of the returned string is reduced to reflect only those characters before the null.

DataOffset is incremented and DataLength is decremented by the value specified regardless of whether the string contains embedded null characters.

The characters in the message data are assumed to be a string in the code page that is specified by the MQMessage.CharacterSet property. Conversion to ActiveX representation is performed for the application.

#### **Defined in:**

MQMessage class

**Syntax:** *string\$* = MQMessage. **ReadNullTerminatedString**(length&)

#### **Parameters:**

length& *Long*. Length of string field in bytes.

### **ReadShort method:**

Starting with the byte referred to by DataOffset, the ReadShort method reads 2 bytes from the Message Data buffer and returns them as an Integer (signed 2-byte) value.

The method fails if MQMessage.DataLength is less than 2 when it is issued.

DataOffset is incremented by 2 and DataLength is decremented by 2 if the method succeeds.

The 2 characters of message data are assumed to be a binary integer. The encoding is specified by the MQMessage.Encoding property.

#### **Defined in:**

MQMessage class

#### **Syntax:**

*integerv%* = MQMessage **.ReadShort**

### **ReadString method:**

This method reads n bytes from the Message Data buffer starting with the byte referred to by DataOffset and returns it as an ActiveX string.

The method fails if MQMessage.DataLength is less than n when it is issued.

DataOffset is incremented by n and DataLength is decremented by n if the method succeeds.

The n characters of message data are assumed to be a string in the code page that is specified by the MQMessage.CharacterSet property. Conversion to ActiveX representation is performed for the application.

**Defined in:** MQMessage class

**Syntax:** *stringv\$* = MQMessage **.ReadString** (*length&* )

#### **Parameter**

*length&* *Long*. Length of string field in bytes.

### ReadUInt2 method:

Starting with the byte referred to by `DataOffset`, the `ReadUInt2` method reads 2 bytes from the Message Data buffer and returns them as a Long (signed 4-byte) integer value.

The method fails if `MQMessage.DataLength` is less than 2 when it is issued.

`DataOffset` is incremented by 2 and `DataLength` is decremented by 2 if the method succeeds.

The 2 bytes of message data are assumed to be an unsigned binary integer. The encoding is specified by the `MQMessage.Encoding` property.

#### Defined in:

`MQMessage` class

#### Syntax:

```
bigint& = MQMessage .ReadUInt2
```

### ReadUnsignedByte method:

Starting with the byte referred to by `DataOffset`, the `ReadUnsignedByte` method reads 1 byte from the Message Data buffer and returns it as an Integer (signed 2-byte) integer value in the range 0 to 255.

The method fails if `MQMessage.DataLength` is less than 1 when it is issued.

`DataOffset` is incremented by 1 and `DataLength` is decremented by 1 if the method succeeds.

The 1 character of message data is assumed to be an unsigned binary integer.

#### Defined in:

`MQMessage` class

#### Syntax:

```
integer% = MQMessage .ReadUnsignedByte
```

### ReadUTF method:

This method reads a UTF format string from the message, starting from the byte referred to by **`DataOffset`**, and returns the UTF format string as an ActiveX string. The UTF format string being read comprises 2 bytes of data that indicate the length of the string, followed by the UTF character data.

The method fails if **`MQMessage.DataLength`** is less than the string length when it is issued.

**`DataOffset`** is incremented by the string length and **`DataLength`** is decremented by the string length if the method succeeds.

#### Defined in:

`MQMessage` class

#### Syntax:

```
value$ = MQMessage .ReadUTF
```

### ResizeBuffer method:

This method alters the amount of storage currently allocated internally to hold the Message Data buffer. It gives the application some control over the automatic buffer management, in that if the application knows that it is going to deal with a large message, it can ensure that a sufficiently large buffer is allocated. The application does not need to use this call - if it does not, the automatic buffer management code will grow the buffer size to fit.

If you resize the buffer to be smaller than the current `MessageLength`, you risk losing data. If you do lose data, the method returns a `CompletionCode` of `MQCC_WARNING` and a `ReasonCode` of `MQRC_DATA_TRUNCATED`.

If you resize the buffer to be smaller than the value of the `DataOffset` property the:

- `DataOffset` property is changed to point to the end of the new buffer
- `DataLength` property is set to zero
- `MessageLength` property is changed to the new buffer size

#### Defined in:

`MQMessage` class

**Syntax:** `MQMessage.ResizeBuffer ( Length& )`

#### Parameter:

`Length&` Long. Size required in characters.

### Write method:

Writes a sequence of bytes to the message buffer from a byte array at the position referred to by `DataOffset`. If necessary the length of the buffer (`MQMessage.MQMessageLength`) is extended to accommodate the full length of the byte array. `DataOffset` is incremented by the number of bytes written if the method succeeds.

#### Defined in:

`MQMessage` class

#### Syntax:

Call `MQMessage.Write(value)`

#### Parameters:

*data*: a byte array or a variant reference to a byte array



**WriteBoolean method:**

Writes a 1-byte Boolean value at the current position in the message buffer from a 2-byte Boolean value. DataOffset is incremented by one.

**Defined in:**

MQMessage class

**Syntax:**

Call `MQMessage.WriteBoolean(value)`

**Parameter:**

*value*: Boolean (2-bytes). Value to be written.

**WriteByte method:**

This method takes a signed 2-byte integer value and writes it into the Message Data buffer as a 1-byte binary number at the position referred to by DataOffset. It replaces any data already at the position in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by one if the method succeeds.

The value specified should be in the range -128 to 127. If it is not, the method returns with CompletionCode MQCC\_FAILED and ReasonCode MQRC\_WRITE\_VALUE\_ERROR.

**Defined in:** MQMessage class

**Syntax:**

Call `MQMessage.WriteByte(value%)`

**Parameter:** *value%* Integer. Value to be written.

**WriteDecimal2 method:**

Writes a signed 2-byte integer as a 2-byte packed decimal number. DataOffset is incremented by two.

**Defined in:**

MQMessage class

**Syntax:**

Call `MQMessage.WriteDecimal2(value%)`

**Parameter:**

*value%* Integer. Value to be written.

**WriteDecimal4 method:**

Writes a signed 4-byte integer as a 4-byte packed decimal number. DataOffset is incremented by four.

**Defined in:**

MQMessage class

**Syntax:**

Call `MQMessage.WroteDecimal4(value&)`

**Parameter:**

*value&* Long. Value to be written.

**WriteDouble method:**

This method takes a signed 8-byte floating point value and writes it into the Message Data buffer as an 8-byte floating point number starting at the position referred to by DataOffset. It replaces any data already at these positions in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by 8 if the method succeeds.

The method converts to the floating point representation specified by the MQMessage.Encoding property. *Conversion to System/360 format is not supported.*

**Defined in:** MQMessage class

**Syntax:**

Call `MQMessage.WriteDouble(value#)`

**Parameter:**

*value#* Double. Value to be written.

**WriteDouble4 method:**

See "ReadDouble4 method" on page 748 for a description of when ReadDouble4 and WriteDouble4 should be used in place of ReadFloat and WriteFloat.

This method takes a signed 8-byte floating point value and writes it into the Message Data buffer as a 4-byte floating number starting at the position referred to by DataOffset.

DataOffset is incremented by 4 if the method succeeds.

It replaces any data already at these positions in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

The method converts to the floating point representation specified by the MQMessage.Encoding property. *Conversion to System/360 format is not supported.*

**Defined in:** MQMessage class

**Syntax:**

Call `MQMessage.WriteDouble4(value#)`

**Parameter:** *value#* Double. Value to be written.

**WriteFloat method:**

This method takes a signed 4-byte floating point value and writes it into the Message Data buffer as a 4-byte floating point number starting at the character referred to by `DataOffset`. It replaces any data already at these positions in the buffer, and extends the length of the buffer (`MQMessage.MessageLength`) if necessary.

`DataOffset` is incremented by 4 if the method succeeds.

The method converts to the binary representation specified by the `MQMessage.Encoding` property. *Conversion to System/360 format is not supported.*

**Defined in:** `MQMessage` class

**Syntax:**

Call `MQMessage.WriteFloat(value!)`

**Parameter** *value!* Float. Value to be written.

**WriteInt2 method:**

This method is identical to the `WriteShort` method.

**Syntax:**

Call `MQMessage.WriteInt2(value%)`

**Parameter** *value%* Integer. Value to be written.

**WriteInt4 method:**

This method is identical to the `WriteLong` method.

**Syntax:**

Call `MQMessage.WriteInt4(value&)`

**Parameter** *value&* Long. Value to be written.

**WriteLong method:**

This method takes a signed 4-byte integer value and writes it into the Message Data buffer as a 4-byte binary number starting at the byte referred to by `DataOffset`. It replaces any data already at these positions in the buffer, and extends the length of the buffer (`MQMessage.MessageLength`) if necessary.

`DataOffset` is incremented by 4 if the method succeeds.

The method converts to the binary representation specified by the `MQMessage.Encoding` property.

**Defined in:** `MQMessage` class

**Syntax:**

Call `MQMessage.WriteLong(value&)`

**Parameter** *value&* Long. Value to be written.

### **WriteNullTerminatedString method:**

This method performs a normal `WriteString` and pads any remaining bytes up to the specified length with null. If the number of bytes written by the initial write string is equal to the specified length then no nulls are written. If the number of bytes exceeds the specified length then an error (reason code `MQRC_WRITE_VALUE_ERROR`) is set.

`DataOffset` is incremented by the specified length if the method succeeds.

**Defined in:** `MQMessage` class

#### **Syntax:**

Call `MQMessage.WriteNullTerminatedString(value$, length&)`

#### **Parameters:**

`value$` String. Value to be written.

`length&` Long. Length of string field in bytes.

### **WriteShort method:**

This method takes a signed 2-byte integer value and writes it into the Message Data buffer as a 2-byte binary number starting at the byte referred to by `DataOffset`. It replaces any data already at these positions in the buffer, and will extend the length of the buffer (`MQMessage.MessageLength`) if necessary.

`DataOffset` is incremented by 2 if the method succeeds.

The method converts to the binary representation specified by the `MQMessage.Encoding` property.

**Defined in:** `MQMessage` class

#### **Syntax:**

Call `MQMessage.WriteShort(value%)`

**Parameter** `value%` Integer. Value to be written.

### **WriteString method:**

This method takes an ActiveX string and writes it into the Message Data buffer starting at the byte referred to by `DataOffset`. It replaces any data already at these positions in the buffer, and will extend the length of the buffer (`MQMessage.MessageLength`) if necessary.

`DataOffset` is incremented by the length of the string in bytes if the method succeeds.

The method converts characters into the code page specified by the `MQMessage.CharacterSet` property.

**Defined in:** `MQMessage` class

#### **Syntax:**

Call `MQMessage.WriteString(value$)`

**Parameter** `value$` String. Value to be written.

### **WriteUInt2 method:**

This method takes a signed 4-byte integer value and writes it into the Message Data buffer as a 2-byte unsigned binary number starting at the byte referred to by `DataOffset`. It replaces any data already at these positions in the buffer, and extends the length of the buffer (`MQMessage.MessageLength`) if necessary.

`DataOffset` is incremented by 2 if the method succeeds.

The method converts to the binary representation specified by the `MQMessage.Encoding` property. The value specified should be in the range 0 to  $2^{*}16-1$ . If it is not the method returns with `CompletionCode MQCC_FAILED` and `ReasonCode MQRC_WRITE_VALUE_ERROR`.

**Defined in:** `MQMessage` class

#### **Syntax:**

Call `MQMessage.WriteUInt2(value&)`

**Parameter** *value&* Long. Value to be written.

### **WriteUnsignedByte method:**

This method takes a signed 2-byte integer value and writes it into the Message Data buffer as a 1-byte unsigned binary number starting at the character referred to by `DataOffset`. It replaces any data already at these positions in the buffer, and extends the length of the buffer (`MQMessage.MessageLength`) if necessary.

`DataOffset` is incremented by 1 if the method succeeds.

The value specified should be in the range 0 to 255. If it is not the method returns with `CompletionCode MQCC_FAILED` and `ReasonCode MQRC_WRITE_VALUE_ERROR`.

#### **Defined in:**

`MQMessage` class

#### **Syntax:**

Call `MQMessage.WriteUnsignedByte(value%)`

**Parameter** *value%* Integer. Value to be written.

### **WriteUTF method:**

This method takes an ActiveX string and writes it into the message data buffer at the current position in UTF format. The data written consists of a 2-byte length followed by the character data. `DataOffset` is incremented by the length of the string if the method succeeds.

#### **Defined in:**

`MQMessage` class

**Syntax:** Call `MQMessage.WriteUTF (value$)`

#### **Parameter:**

*value\$* String. Value to be written.

## MQPutMessageOptions class

This class encapsulates the various options that control the action of putting a message onto an IBM MQ Queue.

### Containment

The MQPutMessageOptions class is contained by the MQSession class.

### Creation

**New** creates a new MQPutMessageOptions object and sets all its properties to initial values.

Alternatively, use the AccessPutMessageOptions method of the MQSession class.

### Syntax

**Dim** *pmo* **As New** MQPutMessageOptions **or**

**Set** *pmo* = **New** MQPutMessageOptions

### Properties

- "CompletionCode property."
- "Options property" on page 759.
- "ReasonCode property" on page 759.
- "ReasonName property" on page 759.
- "RecordFields property" on page 760.
- "ResolvedQueueManagerName property" on page 760.
- "ResolvedQueueName property" on page 760.

### Methods

- "ClearErrorCodes method" on page 760.

### CompletionCode property:

Read-only. Returns the completion code set by the last method or property access issued against the object.

**Defined in:** MQPutMessageOptions class

**Data Type:** Long

### Values:

- MQCC\_OK
- MQCC\_WARNING
- MQCC\_FAILED

**Syntax:** To get: *completioncode* = *PutOpts* .**CompletionCode**

**Options property:**

Read-write. The MQPMO Options field. The initial value of this field is MQPMO\_NONE. For more information, see MQPMO options.

**Defined in:** MQPutMessageOptions Class.

**Data Type:** Long

**Syntax:** To get: *options&* = *PutOpts* .**Options**

To set: *PutOpts* .**Options** = *options&*

The MQPMO\_PASS\_IDENTITY\_CONTEXT and MQPMO\_PASS\_ALL\_CONTEXT options are not supported.

**ReasonCode property:**

Read-only. Returns the reason code set by the last method or property access issued against the object.

**Defined in:** MQPutMessageOptions class

**Data Type:** Long

**Values:**

- See API reason codes.

**Syntax:** To get: *reasoncode&* = *PutOpts* .**ReasonCode**

**ReasonName property:**

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC\_QMGR\_NOT\_AVAILABLE".

**Defined in:** MQPutMessageOptions class

**Data Type:** String

**Values:**

- See API reason codes.

**Syntax:** To get: *reasonname\$* = *PutOpts* .**ReasonName**

**RecordFields property:**

Read-write. Flags indicating which fields are to be customized on a per-queue basis when putting a message to a distribution list. The initial value is zero.

This property corresponds to the PutMsgRecFields flags in the MQI MQPMO structure. In the MQI, these flags control which fields (in the MQPMR structure) are present and used by the MQPUT. In an MQPutMessageOptions object these fields are always present and the flags therefore only affect which fields are used by the Put.

**Defined in:**

MQPutMessageOptions class

**Data Type:**

Long

**Syntax:** To get: *recordfields&* = *PutOpts* .**RecordFields**

To set: *PutOpts* .**RecordFields** = *recordfields&*

**ResolvedQueueManagerName property:**

Read-only. The MQPMO ResolvedQMgrName field. See ResolvedQMgrName (MQCHAR48) for details. The initial value is all blanks.

**Defined in:** MQPutMessageOptions class

**Data Type:** String of 48 characters

**Syntax:** To get: *qmgr\$* = *PutOpts* .**ResolvedQueueManagerName**

**ResolvedQueueName property:**

Read-only. The MQPMO ResolvedQName field. See ResolvedQName (MQCHAR48) for details. The initial value is all blanks.

**Defined in:** MQPutMessageOptions class

**Data Type:** String of 48 characters

**Syntax:** To get: *qname\$* = *PutOpts* .**ResolvedQueueName**

**ClearErrorCodes method:**

Resets the CompletionCode to MQCC\_OK and the ReasonCode to MQRC\_NONE for both the MQPutMessageOptions class and the MQSession class.

**Defined in:**

MQPutMessageOptions class

**Syntax:**

Call *PutOpts* .**ClearErrorCodes()**



## MQGetMessageOptions class

This class encapsulates the various options that control the action of getting a message from an IBM MQ queue.

### Containment

The MQGetMessageOptions class is contained by the MQSession class.

### Creation

**New** creates a new MQGetMessageOptions object and sets all its properties to initial values.

Alternatively, use the AccessGetMessageOptions method of the MQSession class.

### Properties

- “CompletionCode property”
- “MatchOptions property” on page 762
- “Options property” on page 762
- “ReasonCode property” on page 762
- “ReasonName property” on page 762
- “ResolvedQueueName property” on page 763
- “WaitInterval property” on page 763

### Methods

- “ClearErrorCodes method” on page 763

### Syntax

**Dim gmo As New MQGetMessageOptions** or

**Set gmo = New MQGetMessageOptions**

#### CompletionCode property:

Read-only. Returns the completion code set by the last method or property access issued against the object.

**Defined in:** MQGetMessageOptions Class.

**Data Type:** Long

#### Values:

- MQCC\_OK
- MQCC\_WARNING
- MQCC\_FAILED

**Syntax:** To get: *completioncode*& = *GetOpts* .**CompletionCode**

**MatchOptions property:**

Read-write. Options controlling selection criteria used for MQGET. The initial value is MQMO\_MATCH\_MSG\_ID + MQMO\_MATCH\_CORREL\_ID.

**Defined in:**

MQGetMessageOptions class

**Data Type:**

Long

**Values:**

See MatchOptions (MQLONG).

**Syntax:** To get: *matchoptions&* = *GetOpts*. **MatchOptions**

To set: *GetOpts*. **MatchOptions** = *matchoptions&*

**Options property:**

Read-write. The MQGMO Options field. See Options for details. Initial value is MQGMO\_NO\_WAIT.

**Defined in:** MQGetMessageOptions Class.

**Data Type:** Long

**Syntax:** To get: *options&* = *GetOpts* **.Options** To set: *GetOpts* **.Options** = *options&*

**ReasonCode property:**

Read-only. Returns the reason code set by the last method or property access issued against the object.

**Defined in:** MQGetMessageOptions class

**Data Type:** Long

**Values:**

- See API reason codes.

**Syntax:** To get: *reasoncode&* = *GetOpts* **.ReasonCode**

**ReasonName property:**

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC\_QMGR\_NOT\_AVAILABLE". **Defined in:** MQGetMessageOptions class

**Data Type:** String

**Values:**

- See API reason codes.

**Syntax:** To get: *reasonname\$* = *MQGetMessageOptions* **.ReasonName**

### **ResolvedQueueName property:**

Read-only. The MQGMO ResolvedQName field. See ResolvedQName (MQCHAR48) for details. The initial value is all blanks.

**Defined in:** MQGetMessageOptions class

**Data Type:** String of 48 characters

**Syntax:** To get: *qname\$* = *GetOpts* .ResolvedQueueName

### **WaitInterval property:**

Read/write. The MQGMO WaitInterval field. The maximum time, in milliseconds, that the Get waits for a suitable message to arrive - if wait action has been requested by the Options property. This field has an initial value of 0. For details of MQGMO options, see MQGMO.

**Defined in:** MQGetMessageOptions class

**Data Type:** Long

**Syntax:** To get: *wait&* = *GetOpts* .WaitInterval

To set: *GetOpts* .WaitInterval = *wait&*

### **ClearErrorCodes method:**

Resets the CompletionCode to MQCC\_OK and the ReasonCode to MQRC\_NONE for both the MQGetMessageOptions class and the MQSession class.

**Defined in:**

MQGetMessageOptions class

**Syntax:**

Call *GetOpts* .ClearErrorCodes()

## **MQDistributionList class**

This class encapsulates a collection of queues - local, remote, or alias for output.

### **Creation**

**new** creates a new MQDistributionList object.

Alternatively, use the AddDistributionList method of the MQQueueManager class

### **Properties**

- "AlternateUserId property" on page 764
- "CloseOptions property" on page 764
- "CompletionCode property" on page 765
- "ConnectionReference property" on page 765
- "FirstDistributionListItem property" on page 765
- "IsOpen property" on page 766
- "OpenOptions property" on page 766
- "ReasonCode property" on page 766

- “ReasonName property” on page 767

## Method

- “AddDistributionListItem method” on page 767
- “ClearErrorCodes method” on page 768
- “Close method” on page 768
- “Open method” on page 768
- “Put method” on page 768

## Syntax

**Dim** *distlist*. **A s New MQDistributionList** or **Set** *distlist* = **New MQDistributionList**

### AlternateUserId property:

Read-write. The alternate user ID used to validate access to the list of queues when they are opened.

#### Defined in:

MQDistributionList class

#### Data Type:

String of 12 characters

**Syntax:** To get: *altuser\$* = *MQDistributionList*. **AlternateUserId**

To set: *MQDistributionList*. **AlternateUserId** = *altuser\$*

### CloseOptions property:

Read-write. Options used to control what happens when the distribution list is closed. The initial value is MQCO\_NONE.

#### Defined in:

MQDistributionList class

#### Data Type:

Long

#### Values:

- MQCO\_NONE
- MQCO\_DELETE
- MQCO\_DELETE\_PURGE

**Syntax:** To get: *closeopt&* = *MQDistributionList*. **CloseOptions**

To set: *MQDistributionList*. **CloseOptions** = *closeopt&*

**CompletionCode property:**

Read-only. The completion code set by the last method or property access issued against the object.

**Defined in:**

MQDistributionList class

**Data Type:**

Long

**Values:**

- MQCC\_OK
- MQCC\_WARNING
- MQCC\_FAILED

**Syntax:** To get: *completioncode* = MQDistributionList. **CompletionCode**

**ConnectionReference property:**

Read-write. The queue manager to which the distribution list belongs.

**Defined in:**

MQDistributionList class

**Data Type:**

MQQueueManager

**Syntax:** To get: *set queuemanager* = MQDistributionList. **ConnectionReference**

To set: *set MQDistributionList. ConnectionReference* = *queuemanager*

**FirstDistributionListItem property:**

Read-only. The first distribution list item object associated with the distribution list.

**Defined in:**

MQDistributionList class

**Data Type:**

MQDistributionListItem

**Values:**

**Syntax:** To get: *set distributionlistitem* = MQDistributionList. **FirstDistributionListItem**

**IsOpen property:**

Read-only.

**Defined in:**

MQDistributionList class

**Data Type:**

Boolean

**Values:**

- TRUE (-1)
- FALSE (0)

**Syntax:** To get: *IsOpen* = *MQDistributionList*. **IsOpen**

**OpenOptions property:**

Read-write. Options to be used when the distribution list is opened.

**Defined in:**

MQDistributionList class

**Data Type:**

Long

**Values:**

See MQPMO options.

**Syntax:** To get: *openopt&* = *MQDistributionList*. **OpenOptions**

To set: *MQDistributionList*. **OpenOptions** = *openopt&*

**ReasonCode property:**

Read-only. The reason code set by the last method or property access issued against the object.

**Defined in:**

MQDistributionList class

**Data Type:**

Long

**Values:**

See API reason codes.

**Syntax:** To get: *reasoncode&* = *MQDistributionList*. **ReasonCode**

**ReasonName property:**

Read-only. The symbolic name for the ReasonCode. For example "MQRC\_QMGR\_NOT\_AVAILABLE".

**Defined in:**

MQDistributionList class

**Data Type:**

String

**Values:**

See API reason codes.

**Syntax:** To get: *reasonname\$* = *MQDistributionList*. **ReasonName**

**AddDistributionListItem method:**

Use this method to create a new MQDistributionListItem object and associate it with the distribution list object. The queue name parameter is mandatory.

This method inserts a new distribution list item as the first item in an existing list. Specifically, this method creates the following configuration:

- In the distribution list, it sets the **FirstDistributionListItem** property to point to the new distribution list item.
- In the new distribution list item, it sets the following properties:
  - It sets the **DistributionList** property to point to the distribution list.
  - It sets the **PreviousDistributionListItem** property to null.
  - It sets the **NextDistributionListItem** property to point to the distribution list item that was previously first, or to null if there were no previous items in the list.

You cannot use this method to add a new item when the distribution list is open.

**Defined in:**

MQDistributionList class

**Syntax:** set distributionlistitem = *MQDistributionList* .**AddDistributionListItem** (QName\$, QMgrName\$)

**Parameters:**

*QName\$* String. Name of the IBM MQ queue.

*QMgrName\$* String. Name of the IBM MQ queue manager.

**ClearErrorCodes method:**

Resets the CompletionCode to MQCC\_OK and the ReasonCode to MQRC\_NONE for both the MQDistributionList class and the MQSession class.

**Defined in:**

MQDistributionList class

**Syntax:**

Call *MQDistributionList*.ClearErrorCodes()

**Close method:**

Closes a distribution list using the current value of Close options.

**Defined in:**

MQDistributionList class

**Syntax:**

Call *MQDistributionList*. **Close** ()

**Open method:**

Opens each of the queues specified by the **QueueName** and (where appropriate) **QueueManagerName** properties of the distribution list items associated with the current object using the current value of AlternateUserId.

**Defined in:**

MQDistributionList class

**Syntax:**

Call *MQDistributionList*.Open()

**Put method:**

Places a message on each of the queues identified by the distribution list items associated with the distribution list.

**Defined in:**

MQDistributionList class

**Syntax**

Call MQDistributionList. **Put** (Message, PutMsgOptions&)

**Parameters**

*Message* MQMessage object representing the message to be put.

*PutMsgOptions* MQPutMessageOptions object containing options to control the put operation. If not specified, default PutMessageOptions are used.

This method takes an MQMessage object as a parameter. The following distribution list item properties can be altered as a result of this method:

- CompletionCode
- ReasonCode



- ReasonName
- MessageId
- MessageIdHex
- CorrelationId
- CorrelationIdHex
- GroupId
- GroupIdHex
- Feedback
- AccountingToken
- AccountingTokenHex

## **MQDistributionListItem class**

This class encapsulates the MQOR, MQRR, and MQPMR structures and associates them with an owning distribution list.

### **Creation**

Use the AddDistributionListItem Method of the MQDistributionList class

### **Properties**

#### **Methods**

- “AccountingToken property” on page 770.
- “AccountingTokenHex property” on page 771.
- “CompletionCode property” on page 771.
- “CorrelationId property” on page 771.
- “CorrelationIdHex property” on page 772.
- “DistributionList property” on page 772.
- “Feedback property” on page 772.
- “GroupId property” on page 773.
- “GroupIdHex property” on page 773.
- “MessageId property” on page 773.
- “MessageIdHex property” on page 774.
- “NextDistributionListItem property” on page 774.
- “PreviousDistributionListItem property” on page 774.
- “QueueManagerName property” on page 775.
- “QueueName property” on page 775.
- “ReasonCode property” on page 775.
- “ReasonName property” on page 776.
- “ClearErrorCodes method” on page 776.

**Properties::**

- AccountingToken property
- AccountingTokenHex property
- CompletionCode property
- CorrelationId property
- CorrelationIdHex property
- DistributionList property
- Feedback property
- GroupId property
- GroupIdHex property
- MessageId property
- MessageIdHex property
- NextDistributionListItem property
- PreviousDistributionListItem property
- QueueManagerName property
- QueueName property
- ReasonCode property
- ReasonName property

**Methods::**

- ClearErrorCodes method

**Creation::**

Use the AddDistributionListItem Method of the MQDistributionList class

**AccountingToken property:**

Read-write. The AccountingToken to be included in the MQPMR of a message when put on a queue. Its initial value is all nulls.

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 32 characters

**Syntax:** To get: *accountingtoken\$* = *MQDistributionListItem*. **AccountingToken**

To set: *MQDistributionListItem*. **AccountingToken** = *accountingtoken\$*

**AccountingTokenHex property:**

Read-write. The AccountingToken to be included in the MQPMR of a message when put on a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 64 valid hexadecimal characters.

Its initial value is "0...0".

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 64 hexadecimal characters representing 32 ASCII characters.

**Syntax:** To get: *accountingtokenh\$* = MQDistributionListItem. **AccountingTokenHex**

To set: MQDistributionListItem. **AccountingTokenHex** = *accountingtokenh\$*

**CompletionCode property:**

Read-only. The completion code set by the last open or put request issued against the owning distribution list object.

**Defined in:**

MQDistributionListItem class

**Data Type:**

Long

**Values:**

- MQCC\_OK
- MQCC\_WARNING
- MQCC\_FAILED

**Syntax:** To get: *completioncode\$* = MQDistributionListItem. **CompletionCode**

**CorrelationId property:**

Read-write. The CorrelId to be included in the MQPMR of a message when put on a queue. Its initial value is all nulls.

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 24 characters

**Syntax:** To get: *correlid\$* = MQDistributionListItem. **CorrelationId**

To set: MQDistributionListItem. **CorrelationId** = *correlid\$*

**CorrelationIdHex property:**

Read-write. The CorrelId to be included in the MQPMR of a message when put on a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0..0".

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 48 hexadecimal characters representing 24 ASCII characters.

**Syntax:** To get: *correlidh\$ = MQDistributionListItem*. **CorrelationIdHex**

To set: *MQDistributionListItem*. **CorrelationIdHex** = *correlidh\$*

**DistributionList property:**

Read-only. The distribution list with which this distribution list item is associated.

**Defined in:**

MQDistributionListItem class

**Data Type:**

MQDistributionList

**Syntax:** To get: *set distributionlist = MQDistributionListItem*. **DistributionList**

**Feedback property:**

Read-write. The Feedback value to be included in the MQPMR of a message when put on a queue.

**Defined in:**

MQDistributionListItem class

**Data Type:**

Long

**Values:**

See Feedback (MQLONG).

**Syntax:** To get: *feedback& = MQDistributionListItem*. **Feedback**

To set: *MQDistributionListItem*. **Feedback** = *feedback&*

**GroupId property:**

Read-write. The GroupId to be included in the MQPMR of a message when put on a queue. Its initial value is all nulls.

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 24 characters

**Syntax:** To get: *groupid\$* = *MQDistributionListItem*. **GroupId**

To set: *MQDistributionListItem*. **GroupId** = *groupid\$*

**GroupIdHex property:**

Read-write. The GroupId to be included in the MQPMR of a message when put on a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0..0".

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 48 hexadecimal characters representing 24 ASCII characters.

**Syntax:** To get: *groupidh\$* = *MQDistributionListItem*. **GroupIdHex**

To set: *MQDistributionListItem*. **GroupIdHex** = *groupidh\$*

**MessageId property:**

Read-write. The MessageId to be included in the MQPMR of a message when put on a queue. Its initial value is all nulls.

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 24 characters

**Syntax:** To get: *messageid\$* = *MQDistributionListItem*. **MessageId**

To set: *MQDistributionListItem*. **MessageId** = *messageid\$*

**MessageIdHex property:**

Read-write. The MessageId to be included in the MQPMR of a message when put on a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0..0".

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 48 hexadecimal characters representing 24 ASCII characters.

**Syntax:** To get: *messageidh\$ = MQDistributionListItem. MessageIdHex*

To set: *MQDistributionListItem. MessageIdHex = messageidh\$*

**NextDistributionListItem property:**

Read-only. The next distribution list item object associated with the same distribution list.

**Defined in:**

MQDistributionListItem class

**Data Type:**

MQDistributionListItem

**Syntax:** To get: *set distributionlistitem = MQDistributionListItem. NextDistributionListItem*

**PreviousDistributionListItem property:**

Read-only. The previous distribution list item object associated with the same distribution list.

**Defined in:**

MQDistributionListItem class

**Data Type:**

MQDistributionListItem

**Syntax:** To get: *set distributionlistitem = MQDistributionListItem. PreviousDistributionListItem*

**QueueManagerName property:**

Read-write. The IBM MQ queue manager name.

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 48 characters.

**Syntax:** To get: *qmname\$* = *MQDistributionListItem.QueueManagerName*

To set: *MQDistributionListItem.QueueManagerName* = *qmname\$*

**QueueName property:**

Read-write. The IBM MQ queue name.

**Defined in:**

MQDistributionListItem class

**Data Type:**

String of 48 characters.

**Syntax:** To get: *qname\$* = *MQDistributionListItem.QueueName*

To set: *MQDistributionListItem.QueueName* = *qname\$*

**ReasonCode property:**

Read-only. The reason code set by the last open or put issued to the owning distribution list object.

**Defined in:**

MQDistributionListItem class

**Data Type:**

Long

**Values:**

See API reason codes.

**Syntax:** To get:

*reasoncode&* = *MQDistributionListItem.ReasonCode*

**ReasonName property:**

Read-only. The symbolic name for the ReasonCode. For example "MQRC\_QMGR\_NOT\_AVAILABLE".

**Defined in:**

MQDistributionListItem class

**Data Type:**

String

**Values:**

See API reason codes.

**Syntax:** To get: *reasonname\$* = *MQDistributionListItem*. **ReasonName**

**ClearErrorCodes method:**

Resets the CompletionCode to MQCC\_OK and the ReasonCode to MQRC\_NONE for both the MQDistributionListItem class and the MQSession class.

**Defined in:**


MQDistributionListItem class

**Syntax:**

Call *MQDistributionListItem*.ClearErrorCodes

## Tracing IBM MQ Automation Classes for ActiveX

Information on the trace facility provided for IBM MQ Automation Classes for ActiveX, common pitfalls and help on how to avoid them.

 From Version 9.0, IBM MQ support for MicrosoftActive X is deprecated. IBM MQ classes for .NET are the recommended replacement technology. For more information, see Developing .NET applications.

This following section explains the trace facility provided and details common pitfalls, with help to avoid them:

- “Controlling trace for IBM MQ Automation Classes for ActiveX” on page 777
- “When your IBM MQ Automation Classes for ActiveX script fails” on page 778
- “Reason codes IBM MQ Automation Classes for ActiveX” on page 778
- “Code level tool” on page 781



## Controlling trace for IBM MQ Automation Classes for ActiveX

The IBM MQ Automation Classes for ActiveX (MQAX) includes a trace facility to help the service organization identify what is happening when you have a problem.

It shows the paths taken when you run your MQAX script. Unless you have a problem, run with tracing disabled to avoid any unnecessary use of system resources.

There are three environment variables that you set to control trace:

- OMQ\_TRACE
- OMQ\_TRACE\_PATH
- OMQ\_TRACE\_LEVEL

**Note:** Specifying *any* value for the **OMQ\_TRACE** variable switches the trace facility on. Even if you set the **OMQ\_TRACE** variable to OFF, trace is still active. To switch off tracing, do not specify a value for **OMQ\_TRACE**.

1. Click **Start**
2. Click **Control Panel**
3. Double-click **System**
4. Click **Advanced**
5. Click **Environment**
6. In the section titled User variables for (username), click **New**
7. Enter the variable name and a valid value in the appropriate fields and click **OK**
8. Click **OK** to close the Environment Variables window
9. Click **OK** to close the System Properties window
10. Close the Control Panel window

When deciding where you want the trace files written, ensure that you have sufficient authority to write to, as well as read from, the disk.

With tracing enabled, it slows down the running of the MQAX, but it does not affect the performance of your ActiveX or IBM MQ environments. When you no longer need a trace file, you can delete it.

You cannot change the status of the OMQ\_TRACE variable while MQAX is running.

### Trace file name and directory

The trace file name takes the form OMQnnnnn.trc, where *nnnnn* is the ID of the ActiveX process running at the time.

Command	Effect
SET OMQ_TRACE_PATH = drive:\directory	Sets the trace directory in which the trace file is written.
SET OMQ_TRACE_PATH =	Removes any existing setting for the trace directory. When the trace directory is not set, the current working directory (when ActiveX is started) is used.
ECHO %OMQ_TRACE_PATH%	Displays the current setting for the trace directory on Windows.

Command	Effect
SET OMQ_TRACE = xxxxxxxx	Switches on tracing. You enable tracing by putting one or more characters after the '=' sign. For example: SET OMQ_TRACE=yes and SET OMQ_TRACE=no. In both of these examples, tracing is enabled. This setting is only effective for a single window/session.
SET OMQ_TRACE= ECHO %OMQ_TRACE%	Switches off tracing. Displays the contents of the environment variable on Windows.
SET	Displays the contents of all the environment variables on Windows.
SET OMQ_TRACE_LEVEL = 9	Sets the trace level to 9. Values greater than 9 do not produce any additional information in the trace file.

## When your IBM MQ Automation Classes for ActiveX script fails

If your IBM MQ Automation Classes for ActiveX script fails, there are a number of sources of information that you can look at.

### First failure symptom report

Independently of the trace facility, for unexpected and internal errors, a First failure symptom report might be produced.

This report is found in a file named `OMQnnnnn.fdc`, where `nnnnn` is the number of the ActiveX process that is running at the time. You find this file in the working directory from which you started ActiveX or in the path specified in the `OMQ_PATH` environment variable.

### Other sources of information

IBM MQ provides various error logs and trace information, depending on the platform involved. See your Windows application event log.

## Reason codes IBM MQ Automation Classes for ActiveX

Reason codes for IBM MQ Automation Classes for ActiveX (MQAX) that can occur in addition to IBM MQ MQI reason codes.

The following reason codes can occur in addition to those documented for the IBM MQ MQI. For other codes, refer to your IBM MQ application event log.

Reason code	Explanation
MQRC_LIBRARY_LOAD_ERROR (6000)	One or more of the IBM MQ libraries could not be loaded. Check that all IBM MQ libraries are in the correct search path on the system you are using. For example, make sure that the directories containing the IBM MQ libraries are in PATH.
MQRC_CLASS_LIBRARY_ERROR (6001)	One of the IBM MQ classlibrary calls returned an unexpected <b>ReasonCode</b> or <b>CompletionCode</b> value. Check the First Failure Symptom Report for details. Take note of the last method/property and class being used and inform IBM Support of the problem.
MQRC_STRING_LENGTH_TOO_BIG (6002)	An attempt has been made to write a UTF format string with a length greater than 65,535 bytes to the message buffer.

Reason code	Explanation
MQRC_WRITE_VALUE_ERROR (6003)	A value is used that is out of range; for example <code>msg.WriteByte (240)</code> .
MQRC_PACKED_DECIMAL_ERROR (6004)	An attempt has been made to read a packed decimal number from the message buffer, but the data at the data pointer is not in a valid packed data format.
MQRC_FLOAT_CONVERSION_ERROR (6005)	An attempt has been made to read a single or double floating point number from the message buffer but the data at the data pointer is not in an appropriate floating point format.
MQRC_REOPEN_EXCL_INPUT_ERROR (6100)	An open object does not have the correct <b>OpenOptions</b> settings and requires one or more additional options. An implicit reopen is required but closure has been prevented because the queue is open for exclusive input and closure would present a window of opportunity for others potentially to gain access to the queue. Set the <b>OpenOptions</b> values explicitly to cover all eventualities so that implicit reopening is not required.
MQRC_REOPEN_INQUIRE_ERROR (6101)	An open object does not have the correct <b>OpenOptions</b> settings and requires one or more additional options. An implicit reopen is required but closure has been prevented because one or more characteristics of the object need to be checked dynamically prior to closure, and the <b>OpenOptions</b> values do not already include the <b>MQ00_INQUIRE</b> option. Set the <b>OpenOptions</b> values explicitly to include the <b>MQ00_INQUIRE</b> option.
MQRC_REOPEN_SAVED_CONTEXT_ERR (6102)	An open object does not have the correct <b>OpenOptions</b> settings and requires one or more additional options. An implicit reopen is required but closure has been prevented because the queue is open with the <b>MQ00_SAVE_ALL_CONTEXT</b> option, and a destructive Get call has been performed previously. This has caused retained state information to be associated with the open queue and this information would be destroyed by closure. Set the <b>OpenOptions</b> values explicitly to cover all eventualities so that implicit reopening is not required.
MQRC_REOPEN_TEMPORARY_Q_ERROR (6103)	An open object does not have the correct <b>OpenOptions</b> settings and requires one or more additional options. An implicit reopen is required, but closure has been prevented because the queue is a local queue of the definition type <b>MQQDT_TEMPORARY_DYNAMIC</b> , which would be destroyed by closure. Set the <b>OpenOptions</b> values explicitly to cover all eventualities so that implicit reopening is not required.
MQRC_ATTRIBUTE_LOCKED (6104)	An attempt has been made to change the value or attribute of an object while that object is open. Certain attributes, such as <b>AlternateUserId</b> , cannot be changed while an object is open.
MQRC_CURSOR_NOT_VALID (6105)	The browse cursor for an open queue has been invalidated since it was last used by an implicit reopen. Set the <b>OpenOptions</b> values explicitly to cover all eventualities so that implicit reopening is not required.
MQRC_ENCODING_ERROR (6106)	The encoding of the next message item needs to be <b>MQENC_NATIVE</b> encoding for reading.
MQRC_STRUCID_ERROR (6107)	The structure of the ID for the next message item, which is derived from the 4 characters beginning at the data pointer, is either missing or is inconsistent with the type of variable into which the item is being read.

Reason code	Explanation
MQRC_NULL_POINTER (6108)	<p>A null pointer has been supplied where a non-null pointer is either required or implied. This might be caused by using explicit declarations for IBM MQ objects that are used from Visual Basic or Excel as parameters to calls. For example:</p> <ul style="list-style-type: none"> <li>• From Visual Basic, <code>dim msg as Object</code> works correctly, whereas <code>dim msg as MqMessage</code> might not work correctly.</li> <li>• From Visual Basic, with a queue defined and set, <code>dim msg as MqMessageq.put msg</code> works correctly, whereas from Excel this command generates an <code>MQRC_NULL_POINTER</code> exception.</li> </ul>
MQRC_NO_CONNECTION_REFERENCE (6109)	<p>The MQQueue object has lost its connection to the MQQueueManager object. This will occur if the queue manager is disconnected. Delete the MQQueue object.</p>
MQRC_NO_BUFFER (6110)	<p>No buffer is available. For an MQMessage object, a buffer cannot be allocated because there is an internal inconsistency in the object state.</p>
MQRC_BINARY_DATA_LENGTH_ERROR (6111)	<p>The length of the binary data is inconsistent with the length of the target attribute. Zero is a correct length for all attributes. 24 is a correct length for a <b>CorrelationId</b> attribute and for a <b>MessageId</b> attribute. 32 is a correct length for an <b>AccountingToken</b> attribute.</p>
MQRC_BUFFER_NOT_AUTOMATIC (6112)	<p>A user-defined and managed buffer cannot be resized. Because message buffers are system managed, this indicates an internal inconsistency.</p>
MQRC_INSUFFICIENT_BUFFER (6113)	<p>There is insufficient buffer space available after the data pointer to accommodate the request. This could be because the buffer cannot be resized.</p>
MQRC_INSUFFICIENT_DATA (6114)	<p>There is insufficient data after the data pointer to accommodate the read request. Reduce the buffer to the correct size and read the data again.</p>
MQRC_DATA_TRUNCATED (6115)	<p>Data has been truncated when copying from one buffer to another. This could be because the target buffer cannot be resized, or because there is a problem addressing one or other buffer, or because a buffer is being downsized with a smaller replacement.</p>
MQRC_ZERO_LENGTH (6116)	<p>A zero length has been supplied where a positive length is either required or implied.</p>
MQRC_NEGATIVE_LENGTH (6117)	<p>A negative length has been supplied where a zero or positive length is required.</p>
MQRC_NEGATIVE_OFFSET (6118)	<p>A negative offset has been supplied where a zero or positive offset is required.</p>
MQRC_INCONSISTENT_FORMAT (6119)	<p>The format of the next message item is inconsistent with the type of variable into which the item is being read.</p>
MQRC_INCONSISTENT_OBJECT_STATE (6120)	<p>There is an inconsistency between this object, which is open, and the referenced MQQueueManager object, which is not connected.</p>
MQRC_CONTEXT_OBJECT_NOT_VALID (6121)	<p>The MQPutMessageOptions context reference does not reference a valid MQQueue object. The object has been previously destroyed.</p>
MQRC_CONTEXT_OPEN_ERROR (6122)	<p>The MQPutMessageOptions context reference refers to an MQQueue object that could not be opened to establish a context. This might be because the MQQueue object has inappropriate open options. Inspect the referenced object reason code to establish the cause.</p>

Reason code	Explanation
MQRC_STRUC_LENGTH_ERROR (6123)	The length of an internal data structure is inconsistent with its content. For an MQRMH header, the length is insufficient to contain the fixed fields and all offset data.
MQRC_NOT_CONNECTED (6124)	A method failed because a required connection to a queue manager is not available, and a connection cannot be established implicitly.
MQRC_NOT_OPEN (6125)	A method failed because an IBM MQ object is not open, and opening cannot be accomplished implicitly.
MQRC_DISTRIBUTION_LIST_EMPTY (6126)	An MQDistributionList failed to open because there are no MQDistributionListItem objects in the distribution list.  Corrective action: Add at least one MQDistributionListItem object to the distribution list.
MQRC_INCONSISTENT_OPEN_OPTIONS (6127)	A method failed because the object is open, and the open options are inconsistent with the required operation.  Corrective action: Open the object with appropriate open options then retry.
MQRC_WRONG_VERSION (6128)	A method failed because a version number specified or encountered is either incorrect or not supported.

## Code level tool

You might be asked by the IBM Service Team which level of code you have installed.

To find this out, run the 'MQAXLEV' utility program.

From the command prompt, change to the directory containing the MQAX200.dll or add the full path length and enter:

```
MQAXLev MQAX200.dll > MQAXLEV.OUT
```

where MQAXLEV.OUT is the name of the output file.

If you do not specify an output file, the detail is displayed on the screen.

An example output file from code level tool is detailed in the following example:

### Example output file from code level tool

```
5639-B43 (C) Copyright IBM Corp. 1996, 1998. ALL RIGHTS RESERVED.
***** Code Level is 5.1 *****
lib/mqole/mqole.cpp, mqole, p000, p000 L981119 1.8 98/08/21
lib/mqlsx/gmqdyn0a.c, mqlsx, p000, p000 L990212 1.6 99/02/11 16:40:24
lib/mqlsx/pc/gmqdyn1p.c, mqlsx, p000, p000 L990212 1.6 99/02/11 16:44:14
lib/mqlsx/xmqcsa.c, mqole, p000, p000 L990216 1.3 99/02/15 13:24:34
lib/mqlsx/xmqfdca.c, mqlsx, p000, p000 L990212 1.3 99/02/11 16:40:35
lib/mqlsx/xmqtrca.c, mqlsx, p000, p000 L990212 1.5 99/02/11 16:12:02
lib/mqlsx/xmqutila.c, mqlsx, p000, p000 L990212 1.3 99/02/11 16:40:40
lib/mqlsx/xmqutl1a.c, mqlsx, p000, p000 L990212 1.4 99/02/11 16:40:30
lib/mqlsx/xmqcnv1a.c, mqlsx, p000, p000 L990212 1.9 99/02/11 16:40:56
lib/mqlsx/xmqmsg.c, mqole, p000, p000 L990219 1.11 99/02/18 12:12:59
```

## ActiveX interface to the MQAI

For a brief overview of COM interfaces and their use in the MQAI, see “Using the Component Object Model Interface (IBM MQ Automation Classes for ActiveX)” on page 681.

The MQAI enables applications to build and send Programmable Command Format (PCF) commands without directly obtaining and formatting the variable length buffers required for PCF. For more information about the MQAI, see The IBM MQ Administration Interface (MQAI). The MQAI ActiveX MQBag class encapsulates the data bags supported by the MQAI in a way that is possible to use in any language that supports the creation of COM objects; for example, Visual Basic, C++, Java, and other ActiveX scripting clients.

The MQAI ActiveX interface is for use with the MQAX classes that provide a COM interface to the MQI. For more information about the MQAX classes, see “Designing MQAX applications that access non-ActiveX applications” on page 682.

The ActiveX interface provides a single class called MQBag. This class is used to create MQAI data bags and its properties and methods are used to create and work with data items within each bag. The MQBag Execute method sends the bag data to an IBM MQ queue manager as a PCF message and collects the replies.

For more information about MQBag class, its properties and methods, see “The MQBag class.”

The PCF message is sent to the queue manager object specified, optionally using specified request and reply queues. Replies are returned in a new MQBag object. The full set of commands and replies is described in Definitions of the Programmable Command Formats. Commands can be sent to any queue manager in the IBM MQ network by selecting the appropriate request and reply queues.

### The MQBag class

The class, MQBag, is used to create MQBag objects as required. When instantiated, the MQBag class returns a new MQBag object reference.

Create a MQBag object in Visual Basic as follows:

```
Dim mqbagg As MQBag
Set mqbagg = New MQBag
```

### MQBag Property

The properties of MQBag objects are explained in the following list:

- “Item property” on page 783.
- “Count property” on page 784.
- “Options property” on page 785.

### MQBag methods

The methods of the MQBag objects are explained in the following list:

- “Add method” on page 785.
- “AddInquiry method” on page 786.
- “Clear method” on page 786.
- “Execute method” on page 786.
- “FromMessage method” on page 787.
- “ItemType method” on page 788.

- "Remove method" on page 788.
- "Selector method" on page 789.
- "ToMessage method" on page 790.
- "Truncate method" on page 790.

## Error Handling

If an error is detected during an operation on an MQBag object, including those errors returned to the bag by an underlying MQAX or MQAI object, an error exception is raised. The MQBag class supports the COM ISupportErrorInfo interface so the following information is available to your error handling routine:

- Error number: composed of the IBM MQ reason code for the error detected and a COM facility code. The facility field, as standard for COM, indicates the area of responsibility for the error. For errors detected by IBM MQ it is always FACILITY\_ITF.
- Error source: identifies the type and version of the object that detected the error. For errors detected during MQBag operations, the error source is always MQBag.MQBag1.
- Error description: the string giving the symbolic name for the IBM MQ reason code.

How you access the error information depends on your scripting language; for example, in Visual Basic the information is returned in the Err object and the IBM MQ reason code is obtained by subtracting the constant vbObjectError from Err.Number.

**ReasonCode = Err.Number - vbObjectError**

If the MQBag Execute method sends a PCF message and a reply is received, the operation is considered successful although the command sent might have failed. In this case, the reply bag itself contains the completion and error reason codes as described in Definitions of the Programmable Command Formats.

## Item property

### Purpose

The Item property represents an item in a bag. It is used to set or inquire about the value of an item. Use of this property corresponds to the following MQAI calls:

- "mqSetString"
- "mqSetInteger"
- "mqInquireInteger"
- "mqInquireString"
- "mqInquireBag"

in the Programmable command formats reference.

### Format

Item (Selector, ItemIndex, Value)

### Parameters

#### Selector (VARIANT) - input

Selector of the item to be set or inquired.

When inquiring about an item, MQSEL\_ANY\_USER\_SELECTOR is the default. When setting an item, MQIA\_LIST or MQCA\_LIST is the default.

If the Selector is not of type long, MQRC\_SELECTOR\_TYPE\_ERROR results.

This parameter is optional.

**ItemIndex (LONG) - input**

This value identifies the occurrence of the item of the specified selector that is to be set or inquired on. MQIND\_NONE is the default.

This parameter is optional.

**Value (VARIANT) - input/output**

The value returned or the value to be set. When inquiring about an item, the return value can be of type long, string, or MQBag. However, when setting an item, the value must be of type long or string; if not, MQRC\_ITEM\_VALUE\_ERROR results.

**Visual Basic Language Invocation**

When inquiring about a value of an item within a bag:

```
Value = mqbag[.Item]([Selector],
[ItemIndex])
```

For MQBag references:

```
Set abag = mqbag[.Item]([Selector].
[ItemIndex])
```

To set the value of an item in a bag:

```
mqbag[.Item]([Selector],
[ItemIndex]) = Value
```

**Count property****Purpose**

The Count property represents the number of data items within a bag. This property corresponds to the MQAI call, "mqCountItems," in the Programmable command formats reference.

**Format**

**Count** (*Selector*, *Value*)

**Parameters****Selector (VARIANT) - input**

Selector of the data items to be included in the count.

MQSEL\_ALL\_USER\_SELECTORS is the default.

If the Selector is not of type long, MQRC\_SELECTOR\_TYPE\_ERROR is returned.

**Value (LONG) - output**

The number of items in the bag included by the *Selector*.

**Visual Basic Language Invocation**

To return the number of items in a bag:

```
ItemCount = mqbag.Count([Selector])
```



## Options property

### Purpose

The Options property sets options for the use of a bag. This property corresponds to the **Options** parameter of the MQAI call, "mqCreateBag," in the Programmable command formats reference.

### Format

Options (*Options*)

### Parameters

#### Options (LONG) - input/output

The bag options.

**Note:** The bag options must be set *before* data items are added to or set within the bag. If the options are changed when the bag is not empty, MQRC\_OPTIONS\_ERROR results. This applies even if the bag is subsequently cleared.

## Visual Basic Language Invocation

When inquiring about the options of an item within a bag:

```
Options = mqbag.Options
```

To set an option of an item in a bag:

```
mqbag.Options = Options
```

## MQBag methods

The methods of the MQBag objects are explained over the following pages.

### Add method:

#### Purpose

The Add method adds a data item to a bag. This method corresponds to the MQAI calls, "mqAddInteger" and "mqAddString," in the Programmable command formats reference.

#### Format

Add (*Value, Selector*)

#### Parameters

##### Value (VARIANT) - input

Integer or string value of the data item.

##### Selector (VARIANT) - input

Selector identifying the item to be added.

Depending on the type of Value, MQIA\_LIST or MQCA\_LIST is the default. If the **Selector** parameter is not of type long, MQRC\_SELECTOR\_TYPE\_ERROR results.

## Visual Basic Language Invocation

To add an item to a bag:

```
mqbag.Add(Value, [Selector])
```

## **AddInquiry method:**

### **Purpose**

The AddInquiry method adds a selector specifying the attribute to be returned when an administration bag is sent to execute an INQUIRE command. This method corresponds to the MQAI call, "mqAddInquiry," in the Programmable command formats reference.

### **Format**

**AddInquiry** (*Inquiry*)

### **Parameters**

#### **Inquiry (LONG) - input**

Selector of the IBM MQ attribute to be returned by the INQUIRE administration command.

### **Visual Basic Language Invocation**

To use the AddInquiry method:

```
mqbag.AddInquiry(Inquiry)
```

## **Clear method:**

### **Purpose**

The Clear method deletes all data items from a bag. This method corresponds to the MQAI call, "mqClearBag," in the Programmable command formats reference.

### **Format**

**Clear**

### **Visual Basic Language Invocation**

To delete all data items from a bag:

```
mqbag.Clear
```

## **Execute method:**

### **Purpose**

The Execute method sends an administration command message to the command server and waits for any reply messages. This method corresponds to the MQAI call, "mqExecute," in the Programmable command formats reference.

### **Format**

*Execute (QueueManager, Command, OptionsBag, RequestQ, ReplyQ, ReplyBag)*

#### Parameters

**QueueManager (MQQueueManager) - input**

The queue manager to which the application is connected.

**Command (LONG) - input**

The command to be executed.

**OptionsBag (MQBag) - input**

The bag containing options that affect the processing of the call.

**RequestQ (MQQueue) - input**

The queue on which the administration command message will be placed.

**ReplyQ (MQQueue) - input**

The queue on which any reply messages are received.

**ReplyBag (MQBag) - output**

A bag reference containing data from reply messages.

#### Visual Basic Language Invocation

To send an administration command message and wait for any reply messages:

```
Set ReplyBag = mqbag.Execute(QueueManager, Command,
[OptionsBag],[RequestQ],[ReplyQ])
```

#### FromMessage method:

#### Purpose

The FromMessage method loads data from a message into a bag. This method corresponds to the MQAI call, "mqBufferToBag," in the Programmable command formats reference.

#### Format

*FromMessage (Message, OptionsBag)*

#### Parameters

**Message (MQMessage) - input**

The message containing the data to be converted.

**OptionsBag (MQBag) - input**

Options to control the processing of the call.

#### Visual Basic Language Invocation

To load data from a message into a bag:

```
mqbag.FromMessage(Message,[OptionsBag])
```

## ItemType method:

### Purpose

The ItemType method returns the type of the value in a specified item in a bag. This method corresponds to the MQAI call, "mqInquireItemInfo," in the Programmable command formats reference.

### Format

**ItemType** (*Selector, ItemIndex, ItemType*)

### Parameters

#### **Selector (VARIANT) - input**

Selector identifying the item to be inquired.

MQSEL\_ANY\_USER\_SELECTOR is the default. If the **Selector** parameter is not of type long, MQRC\_SELECTOR\_TYPE\_ERROR results.

#### **ItemIndex (LONG) - input**

Index of items to be inquired.

MQIND\_NONE is the default.

#### **ItemType (LONG) - output**

Data type of the specified item.

**Note:** Either the **Selector** parameter, **ItemIndex** parameter, or both must be specified. If neither parameter is present, MQRC\_PARAMETER\_MISSING results.

### Visual Basic Language Invocation

To return the type of a value:

```
ItemType = mqbag.ItemType([Selector],
[ItemIndex])
```

## Remove method:

### Purpose

The Remove method deletes an item from a bag. This method corresponds to the MQAI call, "mqDeleteItem," in the Programmable command formats reference.

### Format

**Remove** (*Selector, ItemIndex*)

### Parameters

#### **Selector (VARIANT) - input**

Selector identifying the item to be deleted.

MQSEL\_ANY\_USER\_SELECTOR is the default. If the **Selector** parameter is not of type long, MQRC\_SELECTOR\_TYPE\_ERROR results.

#### **ItemIndex (LONG) - input**

Index of the item to be deleted.

MQIND\_NONE is the default.

**Note:** Either the **Selector** parameter, **ItemIndex** parameter, or both must be specified. If neither parameter is present, MQRC\_PARAMETER\_MISSING results.

### Visual Basic Language Invocation

To delete an item from a bag:

```
mcbag.Remove([Selector],[ItemIndex])
```

#### Selector method:

#### Purpose

The Selector method returns the selector of a specified item within a bag. This method corresponds to the MQAI call, "mqInquireItemInfo," in the Programmable command formats reference.

#### Format

**Selector** (*Selector, ItemIndex, OutSelector*)

#### Parameters

##### **Selector (VARIANT) - input**

Selector identifying the item to be inquired.

MQSEL\_ANY\_USER\_SELECTOR is the default. If the **Selector** parameter is not of type long, MQRC\_SELECTOR\_TYPE\_ERROR results.

##### **ItemIndex (LONG) - input**

Index of the item to be inquired.

MQIND\_NONE is the default.

##### **OutSelector (VARIANT) - output**

Selector of the specified item.

**Note:** Either the **Selector** parameter, **ItemIndex** parameter, or both must be specified. If neither parameter is present, MQRC\_PARAMETER\_MISSING results.

### Visual Basic Language Invocation

To return the selector of an item:

```
OutSelector = mcbag.Selector([Selector],
[ItemIndex])
```

## **ToMessage method:**

### **Purpose**

The ToMessage method returns a reference to an MQMessage object. The reference contains data from a bag. This method corresponds to the MQAI call, "mqBagToBuffer," in the Programmable command formats reference.

### **Format**

**ToMessage** (*OptionsBag, Message*)

### **Parameters**

#### **OptionsBag (MQBag) - input**

A bag containing options that control the processing of the method.

#### **Message (MQMessage) - output**

An MQMessage object reference containing data from the bag.

### **Visual Basic Language Invocation**

To use the ToMessage Method:

```
Set Message = mqbag.ToMessage([OptionsBag])
```

## **Truncate method:**

### **Purpose**

The Truncate method reduces the number of user items in a bag. This method corresponds to the MQAI call, "mqTruncateBag," in the Programmable command formats reference.

### **Format**

**Truncate** (*ItemCount*)

### **Parameters**

#### **ItemCount (LONG) - input**

The number of user items to remain in the bag after truncation has occurred.

### **Visual Basic Language Invocation**

To reduce the number of user items in a bag:

```
mqbag.Truncate(ItemCount)
```

## About the IBM MQ Automation Classes for ActiveX Starter samples

This appendix describes the IBM MQ Automation Classes for ActiveX Starter samples, and explains how to use them.

IBM MQ for Windows provides the following Visual Basic sample programs:

- MQAXTRIV.VBP
- MQAXBSRV.VBP
- MQAXDLST.VBP
- MQAXCLSS.VBP

These samples run on Visual Basic 4 or Visual Basic 5. You will find them in the directory ... \tools\mqax\samples\vb.

In the same directory you will also find samples for Microsoft Excel and html. These are:

- MQAX.XLS
- MQAXTRIV.XLS
- MQAXTRIV.HTM

**Note:** If using Visual Basic 5, you **must** select and install Visual Basic component grid32.ocx.

### What is demonstrated in the samples

The samples demonstrate how to use IBM MQ Automation Classes for ActiveX to:

- Connect to a queue manager
- Access a queue
- Put a message on a queue
- Get a message from a queue

The central part of the Visual Basic sample is shown on the following pages.

“Preparing to run the samples” on page 792 and

“Error handling in the samples” on page 792

### Running the ActiveX Starter samples

Before you run the IBM MQ Automation Classes for ActiveX Starter samples check that you have a default queue manager running and that you have created the required queue definitions. For details of creating and running a queue manager and creating a queue, refer to Administering. The sample uses the queue SYSTEM.DEFAULT.LOCAL.QUEUE which should be defined on any normally set up IBM MQ server.

The different ways of using data bags are as shown in the following list:

- Connect to a queue manager
- Access a queue
- Put a message on a queue
- Get a message from a queue

For information on the MQAX starter samples for Microsoft Basic Version 4 or later, see “Running the MQAXTRIV sample” on page 792

For information on a sample that allow you to browse properties and methods of queue managers and queue objects, see “Starting the MQAXCLSS sample” on page 794

For information on the MQAXDLST sample, “The MQAXDLST sample” on page 794

For information on running the MQAX starter sample for Microsoft Excel 95 or later, MQAXTRIV.XLS, see “Running the MQAXTRIV.XLS sample” on page 795.

For information on running the Bank demonstration with MQAX.XLS, see “Running the Bank demonstration with MQAX.XLS” on page 795

For information on starter sample using an ActiveX compatible WWW browser, see “Starter sample using an ActiveX compatible WWW browser” on page 795

## Preparing to run the samples

To run any of the samples you need one of the following depending on which of the samples you intend to run.

- Microsoft Visual Basic Version 4 (or later)
- Microsoft Excel 95 (or later)
- A web browser

You also need:

- An IBM MQ queue manager running.
- An IBM MQ queue already defined.

## Error handling in the samples

Most of the samples provided in the IBM MQ Automation Classes for ActiveX package exhibit little or no error handling. For more information about error handling, see “Error handling” on page 687.

## Running the MQAXTRIV sample

1. Start the queue manager.
2. In Windows Explorer or File Manager, select the icon for the sample, MQAXTRIV.VBP (Visual Basic Project file) and open the file.  
The Visual Basic program starts and opens the file, MQAXTRIV.VBP.
3. In Visual Basic, press function key 5 (F5) to run the sample.
4. Click anywhere in the window form, MQAX trivial tester.

If everything is working correctly, the window background should change to green. If there is a problem with your setup, the window background should change to red and error information will be displayed.

The following figure shows the central part of the Visual Basic sample.

Option Explicit

```
Private Sub Form_Click()
```

```
'*****
'* This simple example illustrates how to put and get an IBM MQ message to
'* and from an IBM MQ message queue. The data from the message returned by the
'* get is read and compared with that from the original message.
'*****
Dim MQSess As MQSession '* session object
Dim QMgr As MQQueueManager '* queue manager object
Dim Queue As MQQueue '* queue object
Dim PutMsg As MQMessage '* message object for put
Dim GetMsg As MQMessage '* message object for get
```



```

Dim PutOptions As MQPutMessageOptions '* put message options
Dim GetOptions As MQGetMessageOptions '* get message options
Dim PutMsgStr As String '* put message data string
Dim GetMsgStr As String '* get message data string
'*****
'* Handle errors
'*****
On Error GoTo HandleError

'*****
'* Initialize the current position for the form
'*****
CurrentX = 0
CurrentY = 0

'*****
'* Create the MQSession object and access the MQQueueManager and (local) MQQueue
'*****
Set MQSess = New MQSession
Set QMgr = MQSess.AccessQueueManager("")
Set Queue = QMgr.AccessQueue("SYSTEM.DEFAULT.LOCAL.QUEUE", _
 MQOO_OUTPUT Or MQOO_INPUT_AS_Q_DEF)

'*****
'* Create a new MQMessage object for use with put, add some data then create an
'* MQPutMessageOptions object and put the message
'*****
Set PutMsg = MQSess.AccessMessage()
PutMsgStr = "12345678 " & Time
PutMsg.MessageData = PutMsgStr
Set PutOptions = MQSess.AccessPutMessageOptions()
Queue.Put PutMsg, PutOptions

'*****
'* Create a new MQMessage object for use with get, set the MessageId (to that of
'* the message that was put), create an MQGetMessageOptions object and get the
'* message.
'*
'* Note: Setting the MessageId ensures that the get returns the MQMessage
'* that was put earlier.
'*****

Set GetMsg = MQSess.AccessMessage()
GetMsg.MessageId = PutMsg.MessageId
Set GetOptions = MQSess.AccessGetMessageOptions()
Queue.Get GetMsg, GetOptions
'*****
'* Read the data from the message returned by the get, compare it with
'* that from the original message and output a suitable message.
'*****
GetMsgStr = GetMsg.MessageData
Cls
If GetMsgStr = PutMsgStr Then
 BackColor = RGB(127, 255, 127) '* set to green for ok
 Print
 Print "Message data comparison was successful."
 Print "Message data: "" & GetMsgStr & """"
Else
 BackColor = RGB(255, 255, 127) '* set to amber for compare error
 Print "Compare error: "
 Print "The message data returned by the get did not match the " & _
 "input data from the original message that was put."
 Print
 Print "Input message data: "" & PutMsgStr & """"
 Print "Returned message data: "" & GetMsgStr & """"
End If

```

```

Exit Sub
'*****
'* Handle errors
'*****
HandleError:
Dim ErrMsg As String
Dim StrPos As Integer

Cls
BackColor = RGB(255, 0, 0) '* set to red for error
Print "An error occurred as follows:"
Print ""
If MQSess.CompletionCode <> MQCC_OK Then
 ErrMsg = Err.Description
 StrPos = InStr(ErrMsg, " ") '* search for first blank
 If StrPos > 0 Then
 Print Left(ErrMsg, StrPos) '* print offending MQAX object name
 Else
 Print Error(Err) '* print complete error object
 End If
 Print ""
 Print "IBM MQ Completion Code = " & MQSess.CompletionCode
 Print "IBM MQ Reason Code = " & MQSess.ReasonCode
 Print "(" & MQSess.ReasonName & ")"
Else
 Print "Visual Basic error: " & Err
 Print Error(Err)
End If

Exit Sub

End Sub

```

## Starting the MQAXCLSS sample

This sample allows you to browse properties and methods of queue managers and queue objects.

1. Start the queue manager.
2. Open the file, MQAXCLSS.VBP, by double-clicking on the document icon in Windows Explorer or by clicking File - Open from the file menu in Visual Basic.
3. Start the sample.
4. Enter the appropriate queue manager and queue names then click the corresponding buttons.

## The MQAXDLST sample

The Visual Basic MQAXDLST sample demonstrates the use of a distribution list to send the same message to two queues with one put. To run the sample, do the same as for the MQAXCLSS sample.

## MQAX Starter sample for Microsoft Excel 95 or later

This section explains how to run the MQAX starter sample for Microsoft Excel 95 or later, MQAXTRIV.XLS.

### Running the MQAXTRIV.XLS sample:

1. Start the queue manager.
2. In Explorer or File Manager, select the icon for the MQAX sample MQAXTRIV.XLS.
3. Click the button in the spreadsheet.
4. The screen is updated with a success (or failure) message.

### Running the Bank demonstration with MQAX.XLS:

Follow these steps to run the Bank demonstration.

1. Start the queue manager.
2. Run the IBM MQ MQSC command file, BANK.TST. This sets up the necessary IBM MQ queue definitions.  
To find out how to use an MQSC command file, refer to Script (MQSC) Commands.
3. Run MQAXBRSRV.VBP. This sample program is the server simulating a back-end application, and it has to run with Microsoft Excel.
4. Run MQAX.XLS. This sample is the client IBM MQ demonstration.
5. Select a customer from the list.
6. Click **Submit**.

After a short pause (approximately 3 seconds), the fields populate with values and a bar chart is displayed.

## Starter sample using an ActiveX compatible WWW browser

**Note:** To run this sample, you must be running an ActiveX compatible Web browser. Microsoft Internet Explorer (but not Netscape Navigator) is a compatible Web browser.

### Running the HTML sample

This sample demonstrates how you can invoke MQAX from both VBScript and JavaScript.

1. Start the queue manager.
2. Open the file, "MQAXTRIV.HTM", in your ActiveX compatible Web browser.  
You can do this either by double-clicking the file icon in Windows Explorer or you can choose File - Open from the File menu of your ActiveX compatible Web browser.
3. Follow the instructions on the screen.

## Developing C++ applications

IBM MQ provides C++ classes equivalent to IBM MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI.

IBM WebSphere MQ Version 7.0, enhancements to the IBM MQ programming interfaces are not applied to the C++ classes.

IBM MQ C++ provides the following features:

- Automatic initialization of IBM MQ data structures.
- Just-in-time queue manager connection and queue opening.
- Implicit queue closure and queue manager disconnection.
- Dead-letter header transmission and receipt.
- IMS bridge header transmission and receipt.
- Reference message header transmission and receipt.
- Trigger message receipt.
- CICS bridge header transmission and receipt.
- Work header transmission and receipt.
- Client channel definition.

The following Booch class diagrams show that all the classes are broadly parallel to those IBM MQ entities in the procedural MQI (for example using C) that have either handles or data structures. All classes inherit from the `ImqError` class (see `ImqError` C++ class ), which allows an error condition to be associated with each object.

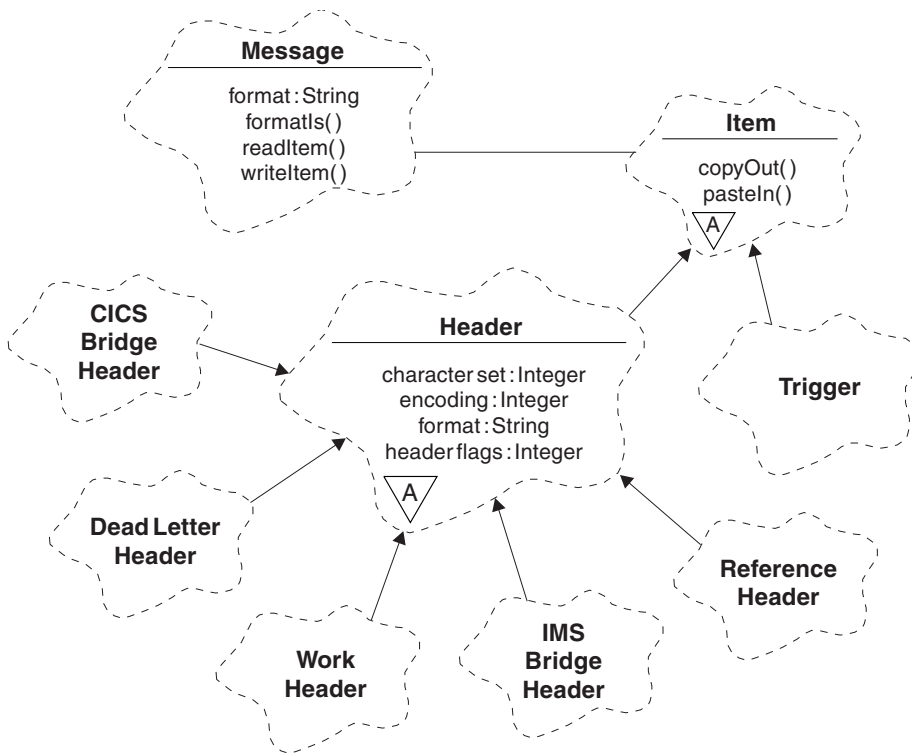


Figure 105. IBM MQ C++ classes (item handling)

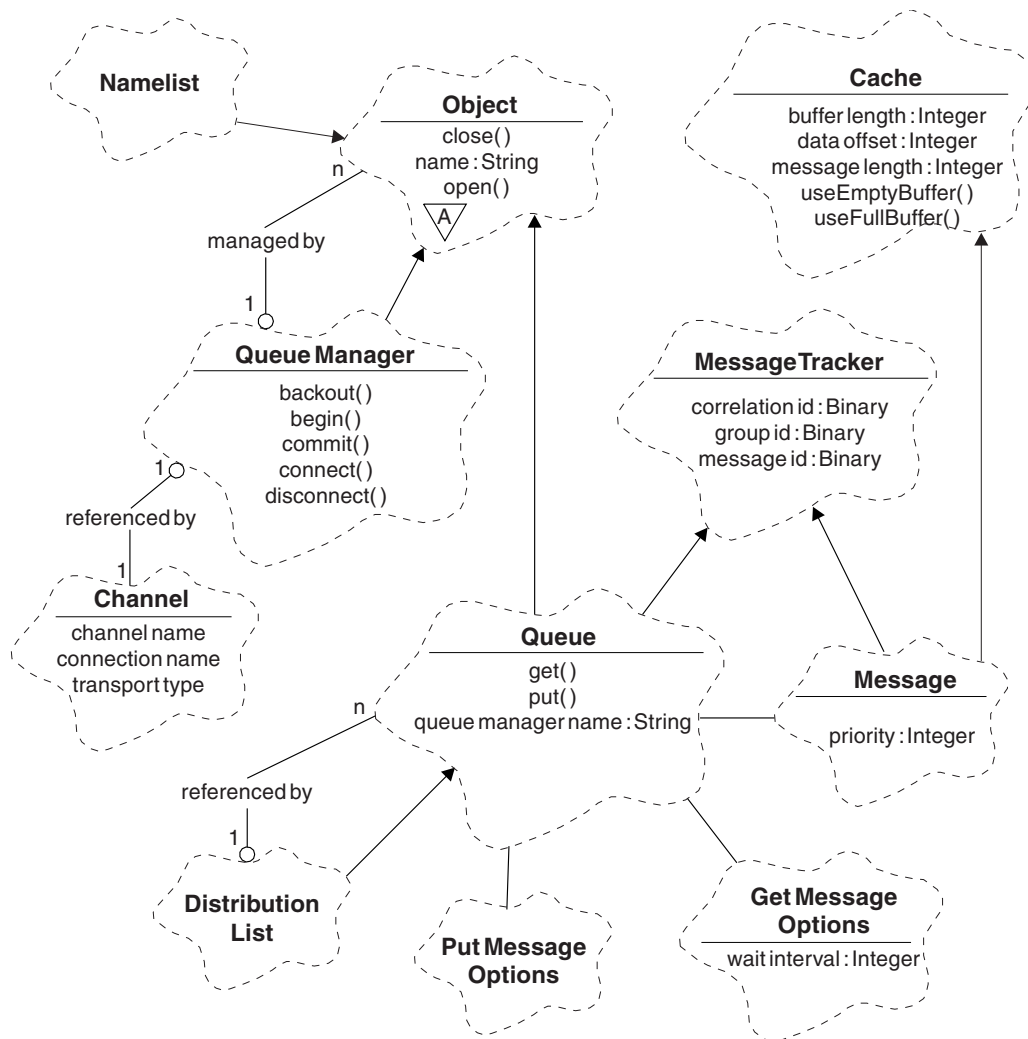


Figure 106. IBM MQ C++ classes (queue management)

To interpret Booch class diagrams correctly, be aware of the following conventions:

- Methods and noteworthy attributes are shown below the *class* name.
- A small triangle within a cloud denotes an *abstract class*.
- *Inheritance* is denoted by an arrow to the parent class.
- An undecorated line between clouds denotes a *cooperative relationship* between classes.
- A line decorated with a number denotes a *referential relationship* between two classes. The number indicates the number of objects that can participate in a particular relationship at any one time.

The following classes and data types are used in the C++ method signatures of the queue management classes (see Figure 106 ) and the item handling classes (see Figure 105 on page 796 ):

- The `ImqBinary` class (see `ImqBinary C++ class` ), which encapsulates byte arrays such as `MQBYTE24`.
- The `ImqBoolean` data type, which is defined as **`typedef unsigned char ImqBoolean`**.
- The `ImqString` class (see `ImqString C++ class` ), which encapsulates character arrays such as `MQCHAR64`.

Entities with data structures are subsumed within appropriate object classes. Individual data structure fields (see C++ and MQI cross-reference ) are accessed with methods.

Entities with handles come under the `ImqObject` class hierarchy (see `ImqObject C++ class`) and provide encapsulated interfaces to the MQI. Objects of these classes exhibit intelligent behavior that can reduce the number of method invocations required relative to the procedural MQI. For example, you can establish and discard queue manager connections as required, or you can open a queue with appropriate options, then close it.

The `ImqMessage` class (see `ImqMessage C++ class`) encapsulates the MQMD data structure and also acts as a holding point for user data and *items* (see “Reading messages in C++” on page 808) by providing cached buffer facilities. You can provide fixed-length buffers for user data and use the buffer many times. The amount of data present in the buffer can vary from one use to the next. Alternatively, the system can provide and manage a buffer of flexible length. Both the size of the buffer (the amount available for receipt of messages) and the amount actually used (either the number of bytes for transmission or the number of bytes actually received) become important considerations.

**Related concepts:**

“C++ sample programs”

Four sample programs are supplied, to demonstrate getting and putting messages.

“C++ language considerations” on page 803

This collection of topics details the aspects of the C++ language usage and conventions that you must consider when writing application programs that use the Message Queue Interface (MQI).

“Preparing message data in C++” on page 807

Message data is prepared in a buffer, which can be supplied by the system or the application. There are advantages to either method. Examples of using a buffer are given.

“Developing applications for IBM MQ” on page 1

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM MQ support applications written in procedural languages, and object oriented languages and frameworks.

**Related reference:**

“Building IBM MQ C++ programs” on page 814

The URL of supported compilers is listed, together with the commands to use to compile, link and run C++ programs and samples on IBM MQ platforms.

**Related information:**

Technical overview

C++ and MQI cross-reference

IBM MQ C++ classes

## C++ sample programs

Four sample programs are supplied, to demonstrate getting and putting messages.



The sample programs are:

- HELLO WORLD (`imqwrlld.cpp`)
- SPUT (`imqspud.cpp`)
- SGET (`imqsget.cpp`)
- DPUT (`imqdput.cpp`)

The sample programs are located in the directories shown in Table 104 on page 799.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Table 104. Location of sample programs

Environment	Directory containing source	Directory containing built programs
AIX	<i>MQ_INSTALLATION_PATH</i> /samp	<i>MQ_INSTALLATION_PATH</i> /samp/bin/ia
 IBM i	/QIBM/ProdData/mqm/samp/	(see note 1 )
HP-UX	<i>MQ_INSTALLATION_PATH</i> /samp	<i>MQ_INSTALLATION_PATH</i> /samp/bin/ah (see note 2 )
 z/OS	thlqual.SCSQCPPS	None
Solaris	<i>MQ_INSTALLATION_PATH</i> /samp	<i>MQ_INSTALLATION_PATH</i> /samp/bin/as
Linux	<i>MQ_INSTALLATION_PATH</i> /samp	<i>MQ_INSTALLATION_PATH</i> /samp/bin/
Windows	<i>MQ_INSTALLATION_PATH</i> \tools\cplusplus\samples	<i>MQ_INSTALLATION_PATH</i> \tools\cplusplus\samples\bin\vn (see note 3 )

#### Notes:

1. Programs built using the ILE C++ compiler for IBM i are in the library QMQM. The include files are in /QIBM/ProdData/mqm/inc.
2. Programs built using the HP ANSI C++ compiler are found in directory *MQ\_INSTALLATION\_PATH*/samp/bin/ah. For further information, see “Building C++ programs on HP-UX” on page 815.
3. Programs built using the Microsoft Visual Studio are found in *MQ\_INSTALLATION\_PATH*\tools\cplusplus\samples\bin\vn. For further information about these compilers, see “Building C++ programs on Windows” on page 820.

### Sample program HELLO WORLD (imqwrld.cpp)

This C++ sample program shows how to put and get a regular datagram (C structure) using the ImqMessage class.

This program shows how to put and get a regular datagram (C structure) using the ImqMessage class. This sample uses few method invocations, taking advantage of implicit method invocations such as **open**, **close**, and **disconnect**.

### On all platforms except z/OS

If you are using a server connection to IBM MQ, follow one of the following procedures:

- To use the existing default queue, SYSTEM.DEFAULT.LOCAL.QUEUE, run the program **imqwrlds** without passing any parameters
- To use a temporary dynamically assigned queue, run **imqwrlds** passing the name of the default model queue, SYSTEM.DEFAULT.MODEL.QUEUE.

If you are using a client connection to IBM MQ, follow one of the following procedures:

- Set up the MQSERVER environment variable (see MQSERVER for more information) and run **imqwrldc**, or
- Run **imqwrldc** passing as parameters the **queue-name**, **queue-manager-name**, and **channel-definition**, where a typical **channel-definition** might be SYSTEM.DEF.SVRCONN/TCP/*hostname* (1414)



## On z/OS

Construct and run a batch job, using the sample JCL `imqwrl.dr`.

See z/OS Batch, RRS Batch and CICS for more information.

### Sample code

```
extern "C" {
#include <stdio.h>
}

#include <imqi.hpp> // IBM MQ C++

#define EXISTING_QUEUE "SYSTEM.DEFAULT.LOCAL.QUEUE"

#define BUFFER_SIZE 12

static char gpszHello[BUFFER_SIZE] = "Hello world" ;
int main (int argc, char * * argv) {
 ImqQueueManager manager ;
 int iReturnCode = 0 ;

 // Connect to the queue manager.
 if (argc > 2) {
 manager.setName(argv[2]);
 }
 if (manager.connect()) {
 ImqQueue * pqueue = new ImqQueue ;
 ImqMessage * pmsg = new ImqMessage ;

 // Identify the queue which will hold the message.
 pqueue -> setConnectionReference(manager);
 if (argc > 1) {
 pqueue -> setName(argv[1]);

 // The named queue can be a model queue, which will result in
 // the creation of a temporary dynamic queue, which will be
 // destroyed as soon as it is closed. Therefore we must ensure
 // that such a queue is not automatically closed and reopened.
 // We do this by setting open options which will avoid the need
 // for closure and reopening.
 pqueue -> setOpenOptions(MQOO_OUTPUT | MQOO_INPUT_SHARED |
 MQOO_INQUIRE);
 } else {
 pqueue -> setName(EXISTING_QUEUE);

 // The existing queue is not a model queue, and will not be
 // destroyed by automatic closure and reopening. Therefore we
 // will let the open options be selected on an as-needed basis.
 // The queue will be opened implicitly with an output option
 // during the "put", and then implicitly closed and reopened
 // with the addition of an input option during the "get".
 }

 // Prepare a message containing the text "Hello world".
 pmsg -> useFullBuffer(gpszHello , BUFFER_SIZE);
 pmsg -> setFormat(MQFMT_STRING);

 // Place the message on the queue, using default put message
 // Options.
 // The queue will be automatically opened with an output option.
 if (pqueue -> put(* pmsg)) {
 ImqString strQueue(pqueue -> name());
 }
 }
}
```



```

// Discover the name of the queue manager.
ImqString strQueueManagerName(manager.name());
printf("The queue manager name is %s.\n",
 (char *)strQueueManagerName);

// Show the name of the queue.
printf("Message sent to %s.\n", (char *)strQueue);

// Retrieve the data message just sent ("Hello world" expected)
// from the queue, using default get message options. The queue
// is automatically closed and reopened with an input option
// if it is not already open with an input option. We get the
// message just sent, rather than any other message on the
// queue, because the "put" will have set the ID of the message
// so, as we are using the same message object, the message ID
// acts as in the message object, a filter which says that we
// are interested in a message only if it has this
// particular ID.

if (pqueue -> get(* pmsg)) {
 int iDataLength = pmsg -> dataLength();

 // Show the text of the received message.
 printf("Message of length %d received, ", iDataLength);

 if (pmsg -> formatIs(MQFMT_STRING)) {
 char * pszText = pmsg -> bufferPointer();

 // If the last character of data is a null, then we can
 // assume that the data can be interpreted as a text
 // string.
 if (! pszText[iDataLength - 1]) {
 printf("text is \"%s\".\n", pszText);
 } else {
 printf("no text.\n");
 }

 } else {
 printf("non-text message.\n");
 }

 } else {
 printf("ImqQueue::get failed with reason code %ld\n",
 pqueue -> reasonCode());
 iReturnCode = (int)pqueue -> reasonCode();
 }

} else {
 printf("ImqQueue::open/put failed with reason code %ld\n",
 pqueue -> reasonCode());
 iReturnCode = (int)pqueue -> reasonCode();
}

// Deletion of the queue will ensure that it is closed.
// If the queue is dynamic then it will also be destroyed.
delete pqueue ;
delete pmsg ;

} else {
 printf("ImqQueueManager::connect failed with reason code %ld\n"
 manager.reasonCode());
 iReturnCode = (int)manager.reasonCode();
}

// Destruction of the queue manager ensures that it is
// disconnected. If the queue object were still available
// and open (which it is not), the queue would be closed

```

```
// prior to disconnection.
return iReturnCode ;
}
```

## Sample programs SPUT (imqspout.cpp) and SGET (imqsget.cpp)

These C++ programs place messages to, and retrieve messages from, a named queue.


These samples show the use of the following classes:

- ImqError (see ImqError C++ class )
- ImqMessage (see ImqMessage C++ class )
- ImqObject (see ImqObject C++ class )
- ImqQueue (see ImqQueue C++ class )
- ImqQueueManager (see ImqQueueManager C++ class )

Follow the appropriate instructions to run the programs.

### On all platforms except z/OS

1. Run **imqspouts** *queue-name*.
2. Type lines of text at the console. These lines are placed as messages onto the specified queue.
3. Enter a null line to end the input.
4. Run **imqsgets** *queue-name* to retrieve all the lines and display them at the console.

 See “Building C++ programs on z/OS Batch, RRS Batch and CICS” on page 821 for more information.

### On z/OS

1. Construct and run a batch job using the sample JCL **imqsputr**. The messages are read from the SYSIN data set.
2. Construct and run a batch job using the sample JCL **imqsgetr**. The messages are retrieved from the queue and sent to the SYSPRINT data set.

## Sample program DPUT (imqdput.cpp)

This C++ sample program puts messages to a distribution list consisting of two queues.

DPUT shows the use of the ImqDistributionList class (see ImqDistributionList C++ class ). This sample is not supported on z/OS.

1. Run **imqdputs** *queue-name-1 queue-name-2* to place messages on the two named queues.
2. Run **imqsgets** *queue-name-1* and **imqsgets** *queue-name-2* to retrieve the messages from those queues.

## C++ language considerations

This collection of topics details the aspects of the C++ language usage and conventions that you must consider when writing application programs that use the Message Queue Interface (MQI).

### C++ Header files

Header files are provided as part of the definition of the MQI, to help you write IBM MQ application programs in the C++ language.

These header files are summarized in the following table.

Table 105. C/C++ header files

Filename	Contents
IMQI.HPP	C++ MQI Classes (includes CMQC.H and IMQTYPE.H)
IMQTYPE.H	Defines the <b>ImqBoolean</b> data type
CMQC.H	MQI data structures and manifest constants

To improve the portability of applications, code the name of the header file in lowercase on the **#include** preprocessor directive:

```
#include <imqi.hpp> // C++ classes
```

### C++ methods and attributes

Method names are in mixed case. Various considerations apply to parameters and return values. Attributes are accessed using set and get methods as appropriate.

Parameters of methods that are *const* are for input only. Parameters with signatures including a pointer (\*) or a reference (&) are passed by reference. Return values that do not include a pointer or a reference are passed by value; in the case of returned objects, these are new entities that become the responsibility of the caller.

Some method signatures include items that take a default if not specified. Such items are always at the end of signatures and are denoted by an equal sign (=); the value after the equal sign indicates the default value that applies if the item is omitted.

All method names in these classes are mixed case, beginning with lowercase. Each word, except the first within a method name, begins with a capital letter. Abbreviations are not used unless their meaning is widely understood. Abbreviations used include *id* (for identity) and *sync* (for synchronization).

Object attributes are accessed using set and get methods. A set method begins with the word *set* ; a get method has no prefix. If an attribute is *read-only*, there is no set method.

Attributes are initialized to valid states during object construction, and the state of an object is always consistent.

## Data types in C++

All data types are defined by the C **typedef** statement.

The type **ImqBoolean** is defined as **unsigned char** in **IMQTYPE.H** and can have the values **TRUE** and **FALSE**. You can use **ImqBinary** class objects in place of **MQBYTE** arrays, and **ImqString** class objects in place of **char \***. Many methods return objects instead of **char** or **MQBYTE** pointers to ease storage management. All return values become the responsibility of the caller, and, in the case of a returned object, the storage can be disposed of using **delete**.

## Manipulating binary strings in C++

Strings of binary data are declared as objects of the **ImqBinary** class. Objects of this class can be copied, compared, and set using the familiar C operators. Example code is provided.

The following code sample shows operations on a binary string:

```
#include <imqi.hpp> // C++ classes

ImqMessage message ;
ImqBinary id, correlationId ;
MQBYTE24 byteId ;

correlationId.set(byteId, sizeof(byteId)); // Set.
id = message.id(); // Assign.
if (correlationId == id) { // Compare.
...
}
```

## Manipulating character strings in C++

Character data is often returned in **ImqString** class objects which can be cast to **char \*** using a conversion operator. The **ImqString** class contains methods to assist in the processing of character strings.

When character data is accepted or returned using MQI C++ methods, the character data is always null-terminated and can be of any length. However, certain limits are imposed by IBM MQ that might result in information being truncated. To ease storage management, character data is often returned in **ImqString** class objects. These objects can be cast to **char \*** using the conversion operator provided, and used for *read-only* purposes in many situations where a **char \*** is required.

**Note:** The **char \*** conversion result from an **ImqString** class object might be null.

Although C functions can be used on the **char \***, there are special methods of the **ImqString** class that are preferable; **operator length ( )** is the equivalent of **strlen** and **storage ( )** indicates the memory allocated for the character data.

## Initial state of objects in C++

All objects have a consistent initial state reflected by their attributes. The initial values are defined in the class descriptions.

## Using C from C++

When you use C functions from a C++ program, include appropriate headers.

The following example shows `string.h` included in a C++ program:

```
extern "C" {
#include <string.h>
}
```

## C++ notational conventions

This example shows how to invoke methods and declare parameters.

This code sample uses the methods and parameters **ImqBoolean ImqQueue::get ( ImqMessage & msg )**

Declare and use the parameters as follows:

```
ImqQueueManager * pmanager ; // Queue manager
ImqQueue * pqueue ; // Message queue
ImqMessage msg ; // Message
char szBuffer[100] ; // Buffer for message data

pmanager = new ImqQueueManager ;
pqueue = new ImqQueue ;
pqueue -> setName("myreplyq");
pqueue -> setConnectionReference(pmanager);

msg.useEmptyBuffer(szBuffer, sizeof(szBuffer));

if (pqueue -> get(msg)) {
 long lDataLength = msg.dataLength();

 ...
}
```

## Implicit operations in C++

Several operations can occur implicitly, *just in time*, to satisfy the prerequisite conditions for the successful execution of a method. These implicit operations are connect, open, reopen, close, and disconnect. You can control connect and open implicit behavior using class attributes.

### Connect

An `ImqQueueManager` object is connected automatically for any method that results in any call to the MQI (see C++ and MQI cross-reference).

### Open

An `ImqObject` object is opened automatically for any method that results in an MQGET, MQINQ, MQPUT, or MQSET call. Use the **openFor** method to specify one or more relevant **open option** values.

### Reopen

An `ImqObject` is reopened automatically for any method that results in an MQGET, MQINQ, MQPUT, or MQSET call, where the object is already open, but the existing **open options** are not adequate to allow the MQI call to be successful. The object is temporarily closed using a temporary **close options** value of MQCO\_NONE. Use the **openFor** method to add a relevant **open option**.

Reopen can cause problems in specific circumstances:

- A temporary dynamic queue is destroyed when it is closed and can never be reopened.
- A queue opened for exclusive input (either explicitly or by default) might be accessed by others in the window of opportunity during closure and reopening.

- A browse cursor position is lost when a queue is closed. This situation does not prevent closure and reopening, but prevents subsequent use of the cursor until MQGMO\_BROWSE\_FIRST is used again.
- The context of the last message retrieved is lost when a queue is closed.

If any of these circumstances occur or can be foreseen, avoid reopening by explicitly setting adequate **open options** before an object is opened (either explicitly or implicitly).

Setting the **open options** explicitly for complex queue-handling situations results in better performance and avoids the problems associated with the use of reopen.

## Close

An ImqObject is closed automatically at any point where the object state is no longer viable, for example if an ImqObject connection reference is severed, or if an ImqObject object is destroyed.

## Disconnect

An ImqQueueManager is disconnected automatically at any point where the connection is no longer viable, for example if an ImqObject connection reference is severed, or if an ImqQueueManager object is destroyed.

## Binary and character strings in C++

The ImqString class encapsulates the traditional *char \** data format. The ImqBinary class encapsulates the binary byte array. Some methods that set character data might truncate the data.

Methods that set character ( **char \*** ) data always take a copy of the data, but some methods might truncate the copy, because certain limits are imposed by IBM MQ.

The ImqString class (see ImqString C++ class ) encapsulates the traditional **char \*** and provides support for:

- Comparison
- Concatenation
- Copying
- Integer-to-text and text-to-integer conversion
- Token (word) extraction
- Uppercase translation

The ImqBinary class (see ImqBinary C++ class ) encapsulates binary byte arrays of arbitrary size. In particular, it is used to hold the following attributes:

- **accounting token** (MQBYTE32)
- **connection tag** (MQBYTE128)
- **correlation id** (MQBYTE24)
- **facility token** (MQBYTE8)
- **group id** (MQBYTE24)
- **instance id** (MQBYTE24)
- **message id** (MQBYTE24)
- **message token** (MQBYTE16)
- **transaction instance id** (MQBYTE16)

Where these attributes belong to objects of the following classes:

- ImqCICSBridgeHeader (see ImqCICSBridgeHeader C++ class )
- ImqGetMessageOptions (see ImqGetMessageOptions C++ class )

- ImqIMSBridgeHeader (see ImqIMSBridgeHeader C++ class )
- ImqMessageTracker (see ImqMessageTracker C++ class )
- ImqQueueManager (see ImqQueueManager C++ class )
- ImqReferenceHeader (see ImqReferenceHeader C++ class )
- ImqWorkHeader (see ImqWorkHeader C++ class )

The ImqBinary class also provides support for comparison and copying.

## Unsupported functions in C++

The IBM MQ C++ classes and methods are independent of IBM MQ platform. They might therefore offer some functions that are not supported on certain platforms.

If you try to use a function on a platform on which it is not supported, the function is detected by IBM MQ but not by the C++ language bindings. IBM MQ reports the error to your program, like any other MQI error.

## Messaging in C++

This collection of topics details how to prepare, read, and write messages in C++.

### Preparing message data in C++

Message data is prepared in a buffer, which can be supplied by the system or the application. There are advantages to either method. Examples of using a buffer are given.

When you send a message, message data is first prepared in a buffer managed by an ImqCache object (see ImqCache C++ class ). A buffer is associated (by inheritance) with each ImqMessage object (see ImqMessage C++ class ): it can be supplied by the application (using either the **useEmptyBuffer** or **useFullBuffer** method) or automatically by the system. The advantage of the application supplying the message buffer is that no data copying is necessary in many cases because the application can use prepared data areas directly. The disadvantage is that the supplied buffer is of a fixed length.

The buffer can be reused, and the number of bytes transmitted can be varied each time, by using the **setMessageLength** method before transmission.

When supplied automatically by the system, the number of bytes available is managed by the system, and data can be copied into the message buffer using, for example, the ImqCache **write** method, or the ImqMessage **writeItem** method. The message buffer grows according to need. As the buffer grows, there is no loss of previously written data. A large or multipart message can be written in sequential pieces.

The following examples show simplified message sends.

1. Use prepared data in a user-supplied buffer

```
char szBuffer[] = "Hello world" ;

msg.useFullBuffer(szBuffer, sizeof(szBuffer));
msg.setFormat(MQFMT_STRING);
```

2. Use prepared data in a user-supplied buffer, where the buffer size exceeds the data size

```
char szBuffer[24] = "Hello world" ;

msg.useEmptyBuffer(szBuffer, sizeof(szBuffer));
msg.setFormat(MQFMT_STRING);
msg.setMessageLength(12);
```

3. Copy data to a user-supplied buffer

```
char szBuffer[12];

msg.useEmptyBuffer(szBuffer, sizeof(szBuffer));
msg.setFormat(MQFMT_STRING);
msg.write(12, "Hello world");
```

#### 4. Copy data to a system-supplied buffer

```
msg.setFormat(MQFMT_STRING);
msg.write(12, "Hello world");
```

#### 5. Copy data to a system-supplied buffer using objects (objects set the message format as well as content)

```
ImqString strText("Hello world");

msg.writeItem(strText);
```

## Reading messages in C++

A buffer can be supplied by the application or the system. Data can be accessed directly from the buffer or read sequentially. There is a class equivalent to each message type. Sample code is given.

When receiving data, the application or the system can supply a suitable message buffer. The same buffer can be used for both multiple transmission and multiple receipt for a particular `ImqMessage` object. If the message buffer is supplied automatically, it grows to accommodate whatever length of data is received. However, a message buffer supplied by the application might not be big enough to hold the data received. Then either truncation or failure might occur, depending on the options used for message receipt.

Incoming data can be accessed directly from the message buffer, in which case the data length indicates the total amount of incoming data. Alternatively, incoming data can be read sequentially from the message buffer. In this case, the data pointer addresses the next byte of incoming data, and the data pointer and data length are updated each time data is read.

*Items* are pieces of a message, all in the user area of the message buffer, that need to be processed sequentially and separately. Apart from regular user data, an item might be a dead-letter header or a trigger message. Items are always associated with message formats; message formats are **not** always associated with items.

There is a class of object for each item that corresponds to a recognizable IBM MQ message format. There is one for a dead-letter header and one for a trigger message. There is no class of object for user data. That is, once the recognizable formats have been exhausted, processing the remainder is left to the application program. Classes for user data can be written by specializing the `ImqItem` class.

The following example shows a message receipt that takes account of a number of potential items that can precede the user data, in an imaginary situation. Non-item user data is defined as anything that occurs after items that can be identified. An automatic buffer (the default) is used to hold an arbitrary amount of message data.

```
ImqQueue queue ;
ImqMessage msg ;

if (queue.get(msg)) {

 /* Process all items of data in the message buffer. */
 do while (msg.dataLength()) {
 ImqBoolean bFormatKnown = FALSE ;
 /* There remains unprocessed data in the message buffer. */

 /* Determine what kind of item is next. */

 if (msg.formatIs(MQFMT_DEAD_LETTER_HEADER)) {
 ImqDeadLetterHeader header ;
 /* The next item is a dead-letter header. */
 /* For the next statement to work and return TRUE, */
 /* the correct class of object pointer must be supplied. */
 bFormatKnown = TRUE ;

 if (msg.readItem(header)) {
```



```

 /* The dead-letter header has been extricated from the */
 /* buffer and transformed into a dead-letter object. */
 /* The encoding and character set of the dead-letter */
 /* object itself are MQENC_NATIVE and MQCCSI_Q_MGR. */
 /* The encoding and character set from the dead-letter */
 /* header have been copied to the message attributes */
 /* to reflect any remaining data in the buffer. */
 /*
 /* Process the information in the dead-letter object. */
 /* Note that the encoding and character set have */
 /* already been processed. */
 ...
}
/* There might be another item after this, */
/* or just the user data. */
}
if (msg.formatIs(MQFMT_TRIGGER)) {
 ImqTrigger trigger ;
 /* The next item is a trigger message. */
 /* For the next statement to work and return TRUE, */
 /* the correct class of object pointer must be supplied. */
 bFormatKnown = TRUE ;
 if (msg.readItem(trigger)) {

 /* The trigger message has been extricated from the */
 /* buffer and transformed into a trigger object. */
 /* Process the information in the trigger object. */
 ...
 }

 /* There is usually nothing after a trigger message. */
}

if (msg.formatIs(FMT_USERCLASS)) {
 UClass object ;
 /* The next item is an item of a user-defined class. */
 /* For the next statement to work and return TRUE, */
 /* the correct class of object pointer must be supplied. */
 bFormatKnown = TRUE ;

 if (msg.readItem(object)) {
 /* The user-defined data has been extricated from the */
 /* buffer and transformed into a user-defined object. */

 /* Process the information in the user-defined object. */
 ...
 }

 /* Continue looking for further items. */
}
if (! bFormatKnown) {
 /* There remains data that is not associated with a specific*/
 /* item class. */
 char * pszDataPointer = msg.dataPointer(); /* Address.*/
 int iDataLength = msg.dataLength(); /* Length. */

 /* The encoding and character set for the remaining data are */
 /* reflected in the attributes of the message object, even */
 /* if a dead-letter header was present. */
 ...
}
}
}

```

In this example, `FMT_USERCLASS` is a constant representing the 8-character format name associated with an object of class `UserClass`, and is defined by the application.

`UserClass` is derived from the `ImqItem` class (see `ImqItem C++ class`), and implements the virtual `copyOut` and `pasteIn` methods from that class.

The next two examples show code from the `ImqDeadLetterHeader` class (see `ImqDeadLetterHeader C++ class`). The first example shows custom-encapsulated message- *writing* code.

```
// Insert a dead-letter header.
// Return TRUE if successful.
ImqBoolean ImqDeadLetterHeader :: copyOut (ImqMessage & msg) {
 ImqBoolean bSuccess ;
 if (msg.moreBytes(sizeof(omqdlh))) {
 ImqCache cacheData(msg); // Preserve original message content.
 // Note original message attributes in the dead-letter header.
 setEncoding(msg.encoding());
 setCharacterSet(msg.characterSet());
 setFormat(msg.format());

 // Set the message attributes to reflect the dead-letter header.
 msg.setEncoding(MQENC_NATIVE);
 msg.setCharacterSet(MQCCSI_Q_MGR);
 msg.setFormat(MQFMT_DEAD_LETTER_HEADER);
 // Replace the existing data with the dead-letter header.
 msg.clearMessage();
 if (msg.write(sizeof(omqdlh), (char *) & omqdlh)) {
 // Append the original message data.
 bSuccess = msg.write(cacheData.messageLength(),
 cacheData.bufferPointer());
 } else {
 bSuccess = FALSE ;
 }
 } else {
 bSuccess = FALSE ;
 }
 // Reflect and cache error in this object.
 if (! bSuccess) {
 setReasonCode(msg.reasonCode());
 setCompletionCode(msg.completionCode());
 }

 return bSuccess ;
}
```

The second example shows custom-encapsulated message- *reading* code.

```
// Read a dead-letter header.
// Return TRUE if successful.
ImqBoolean ImqDeadLetterHeader :: pasteIn (ImqMessage & msg) {
 ImqBoolean bSuccess = FALSE ;

 // First check that the eye-catcher is correct.
 // This is also our guarantee that the "character set" is correct.
 if (ImqItem::structureIdIs(MQDLH_STRUC_ID, msg)) {
 // Next check that the "encoding" is correct, as the MQDLH
 // contains numeric data.
 if (msg.encoding() == MQENC_NATIVE) {

 // Finally check that the "format" is correct.
 if (msg.formatIs(MQFMT_DEAD_LETTER_HEADER)) {
 char * pszBuffer = (char *) & omqdlh ;
 // Transfer the MQDLH from the message and move pointer on.
 if (bSuccess = msg.read(sizeof(omdlh), pszBuffer)) {
 // Update the encoding, character set and format of the
 // message to reflect the remaining data.
 }
 }
 }
 }
}
```

```

 msg.setEncoding(encoding());
 msg.setCharacterSet(characterSet());
 msg.setFormat(format());
 } else {

 // Reflect the cache error in this object.
 setReasonCode(msg.reasonCode());
 setCompletionCode(msg.completionCode());
 }
 } else {
 setReasonCode(MQRC_INCONSISTENT_FORMAT);
 setCompletionCode(MQCC_FAILED);
 }
 } else {
 setReasonCode(MQRC_ENCODING_ERROR);
 setCompletionCode(MQCC_FAILED);
 }
 {
 } else {
 setReasonCode(MQRC_STRUC_ID_ERROR);
 setCompletionCode(MQCC_FAILED);
 }
 }

 return bSuccess ;
}

```

With an automatic buffer, the buffer storage is *volatile*. That is, buffer data might be held at a different physical location after each **get** method invocation. Therefore, each time buffer data is referenced, use the **bufferPointer** or **dataPointer** methods to access message data.

You might want a program to set aside a fixed area for receiving message data. In this case, invoke the **useEmptyBuffer** method before using the **get** method.

Using a fixed, nonautomatic area limits messages to a maximum size, so it is important to consider the `MQGMO_ACCEPT_TRUNCATED_MSG` option of the `ImqGetMessageOptions` object. If this option is not specified (the default), the `MQRC_TRUNCATED_MSG_FAILED` reason code can be expected. If this option is specified, the `MQRC_TRUNCATED_MSG_ACCEPTED` reason code might be expected depending on the design of the application.

The next example shows how a fixed area of storage can be used to receive messages:

```

char * pszBuffer = new char[100];

msg.useEmptyBuffer(pszBuffer, 100);
gmo.setOptions(MQGMO_ACCEPT_TRUNCATED_MSG);
queue.get(msg, gmo);

delete [] pszBuffer ;

```

In this code fragment, the buffer can always be addressed directly, with *pszBuffer*, as opposed to using the **bufferPointer** method. However, it is better to use the **dataPointer** method for general-purpose access. The application (not the `ImqCache` class object) must discard a user-defined (nonautomatic) buffer.

**Attention:** Specifying a null pointer and zero length with **useEmptyBuffer** does not nominate a fixed-length buffer of length zero as might be expected. This combination is interpreted as a request to ignore any previous user-defined buffer, and instead revert to the use of an automatic buffer.

## Writing a message to the dead-letter queue in C++

Example program code for writing a message to the dead-letter queue.

A typical case of a multipart message is one containing a dead-letter header. The data from a message that cannot be processed is appended to the dead-letter header.

```
ImqQueueManager mgr ; // The queue manager.
ImqQueue queueIn ; // Incoming message queue.
ImqQueue queueDead ; // Dead-letter message queue.
ImqMessage msg ; // Incoming and outgoing message.
ImqDeadLetterHeader header ; // Dead-letter header information.

// Retrieve the message to be rerouted.
queueIn.setConnectionReference(mgr);
queueIn.setName(MY_QUEUE);
queueIn.get(msg);

// Set up the dead-letter header information.
header.setDestinationQueueManagerName(mgr.name());
header.setDestinationQueueName(queueIn.name());
header.setPutApplicationName(/* ? */);
header.setPutApplicationType(/* ? */);
header.setPutDate(/* TODAY */);
header.setPutTime(/* NOW */);
header.setDeadLetterReasonCode(FB_APPL_ERROR_1234);

// Insert the dead-letter header information. This will vary
// the encoding, character set and format of the message.
// Message data is moved along, past the header.
msg.writeItem(header);

// Send the message to the dead-letter queue.
queueDead.setConnectionReference(mgr);
queueDead.setName(mgr.deadLetterQueueName());
queueDead.put(msg);
```

## Writing a message to the IMS bridge in C++

Example program code for writing a message to the IMS bridge.

Messages sent to the IBM MQ - IMS bridge might use a special header. The IMS bridge header is prefixed to regular message data.

```
ImqQueueManager mgr; // The queue manager.
ImqQueue queueBridge; // IMS bridge message queue.
ImqMessage msg; // Outgoing message.
ImqIMSBridgeHeader header; // IMS bridge header.

// Set up the message.
//
// Here we are constructing a message with format
// MQFMT_IMS_VAR_STRING, and appropriate data.
//
msg.write(2, /* ? */); // Total message length.
msg.write(2, /* ? */); // IMS flags.
msg.write(7, /* ? */); // Transaction code.
msg.write(/* ? */ , /* ? */); // String data.
msg.setFormat(MQFMT_IMS_VAR_STRING); // The format attribute.

// Set up the IMS bridge header information.
//
// The reply-to-format is often specified.
// Other attributes can be specified, but all have default values.
//
header.setReplyToFormat(/* ? */);

// Insert the IMS bridge header into the message.
```

```

//
// This will:
// 1) Insert the header into the message buffer, before the existing
// data.
// 2) Copy attributes out of the message descriptor into the header,
// for example the IMS bridge header format attribute will now
// be set to MQFMT_IMS_VAR_STRING.
// 3) Set up the message attributes to describe the header, in
// particular setting the message format to MQFMT_IMS.
//
msg.writeItem(header);

// Send the message to the IMS bridge queue.
//
queueBridge.setConnectionReference(mgr);
queueBridge.setName(/* ? */);
queueBridge.put(msg);

```

## Writing a message to the CICS bridge in C++

Example program code for writing a message to the CICS bridge.

Messages sent to IBM MQ for z/OS using the CICS bridge require a special header. The CICS bridge header is prefixed to regular message data.

```

ImqQueueManager mgr ; // The queue manager.
ImqQueue queueIn ; // Incoming message queue.
ImqQueue queueBridge ; // CICS bridge message queue.
ImqMessage msg ; // Incoming and outgoing message.
ImqCicsBridgeHeader header ; // CICS bridge header information.

// Retrieve the message to be forwarded.
queueIn.setConnectionReference(mgr);
queueIn.setName(MY_QUEUE);
queueIn.get(msg);

// Set up the CICS bridge header information.
// The reply-to format is often specified.
// Other attributes can be specified, but all have default values.
header.setReplyToFormat(/* ? */);

// Insert the CICS bridge header information. This will vary
// the encoding, character set and format of the message.
// Message data is moved along, past the header.
msg.writeItem(header);

// Send the message to the CICS bridge queue.
queueBridge.setConnectionReference(mgr);
queueBridge.setName(/* ? */);
queueBridge.put(msg);

```

## Writing a message with a work header in C++

Example program code for writing a message destined for a queue managed by the z/OS Workload Manager.

Messages sent to IBM MQ for z/OS, which are destined for a queue managed by the z/OS Workload Manager, require a special header. The work header is prefixed to regular message data.

```
ImqQueueManager mgr ; // The queue manager.
ImqQueue queueIn ; // Incoming message queue.
ImqQueue queueWLM ; // WLM managed queue.
ImqMessage msg ; // Incoming and outgoing message.
ImqWorkHeader header ; // Work header information

// Retrieve the message to be forwarded.
queueIn.setConnectionReference(mgr);
queueIn.setName(MY_QUEUE);
queueIn.get(msg);

// Insert the Work header information. This will vary
// the encoding, character set and format of the message.
// Message data is moved along, past the header.
msg.writeItem(header);

// Send the message to the WLM managed queue.
queueWLM.setConnectionReference(mgr);
queueWLM.setName(/* ? */);
queueWLM.put(msg);
```

## Building IBM MQ C++ programs

The URL of supported compilers is listed, together with the commands to use to compile, link and run C++ programs and samples on IBM MQ platforms.

For a list of the compilers for each supported platform and version of IBM MQ, see System Requirements for IBM MQ.

The command you need to compile and link your IBM MQ C++ program depends on your installation and requirements. The examples that follow show typical compiling and linking commands for some of the compilers using the default installation of IBM MQ on a number of platforms.

## Building C++ programs on AIX

Build IBM MQ C++ programs on AIX using the XL C Enterprise Edition compiler.

### Client

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

#### 32-bit unthreaded application

```
x1C -o imqsputc_32 imqsput.cpp -qchars=signed -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -limqc23ia -limqb23ia -lmqic
```

#### 32-bit threaded application

```
x1C_r -o imqsputc_32_r imqsput.cpp -qchars=signed -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -limqc23ia_r -limqb23ia_r -lmqic_r
```

#### 64-bit unthreaded application

```
x1C -q64 -o imqsputc_64 imqsput.cpp -qchars=signed -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -limqc23ia -limqb23ia -lmqic
```

#### 64-bit threaded application

```
x1C_r -q64 -o imqsputc_64_r imqsput.cpp -qchars=signed -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -limqc23ia_r -limqb23ia_r -lmqic_r
```

## Server

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

### 32-bit unthreaded application

```
x1C -o imqsput_32 imqsput.cpp -qchars=signed -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -limqs23ia -limqb23ia -lmqm
```

### 32-bit threaded application

```
x1C_r -o imqsput_32_r imqsput.cpp -qchars=signed -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -limqs23ia_r -limqb23ia_r -lmqm_r
```

### 64-bit unthreaded application

```
x1C -q64 -o imqsput_64 imqsput.cpp -qchars=signed -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -limqs23ia -limqb23ia -lmqm
```

### 64-bit threaded application

```
x1C_r -q64 -o imqsput_64_r imqsput.cpp -qchars=signed -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -limqs23ia_r -limqb23ia_r -lmqm_r
```

## Building C++ programs on HP-UX

Build IBM MQ C++ programs on HP-UX using the aC++ or aCC compilers.

On HP-UX Itanium, IBM MQ supports only the Standard runtime. Use the aCC compiler.

- libimqi23bh.sl provides the IBM MQ C++ classes for the Standard runtime.
- For compatibility with earlier releases, a symbolic link is provided from libimqi23ah.sl to libimqi23bh.sl.

## IA64 (IPF)

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

### Client: IA64 (IPF)

#### 32-bit unthreaded application

```
aCC -Wl,+b,: +e -D_HPUX_SOURCE -o imqsputc_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -limqi23bh -lmqic
```

#### 32-bit threaded application

```
aCC -Wl,+b,: +e -D_HPUX_SOURCE -o imqsputc_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -limqi23bh_r -lmqic_r -lpthread
```

#### 64-bit unthreaded application

```
aCC +DD64 +e -D_HPUX_SOURCE -o imqsputc_64 imqsput.cpp
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -limqi23bh -lmqic
```

#### 64-bit threaded application

```
aCC +DD64 +e -D_HPUX_SOURCE -o imqsputc_64_r imqsput.cpp
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -limqi23bh_r -lmqic_r
-lpthread
```

### Server: IA64 (IPF)

#### 32-bit unthreaded application

```
aCC -Wl,+b,: +e -D_HPUX_SOURCE -o imqsput_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -limqi23bh -lmqm
```

#### 32-bit threaded application

```
aCC -Wl,+b,: +e -D_HPUX_SOURCE -o imqsput_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -limqi23bh_r -lmqm_r -lpthread
```

#### 64-bit unthreaded application

```
aCC +DD64 +e -D_HPUX_SOURCE -o imqsput_64 imqsput.cpp
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -limqi23bh -lmqm
```

## 64-bit threaded application

```
aCC +DD64 +e -D_HPUX_SOURCE -o imqspu64_r imqspu64.cpp
-I MQ_INSTALLATION_PATH/inc -L MQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -limq23bh_r -lmqm_r
-lpthread
```

## Building C++ programs on IBM i

Build IBM MQ C++ programs on IBM i using the ILE C++ compiler.

IBM ILE C++ for IBM i is a native compiler for C++ programs. The following instructions describe how to use this compiler to create IBM MQ C++ applications using the *Hello World!* IBM MQ sample program as an example.

1. Install the ILE C++ for IBM i compiler as directed in the *Read Me first!* manual that accompanies the product.
2. Ensure that the QCXXN library is in your library list.
3. Create the HELLO WORLD sample program:
  - a. Create a module:

```
CRTCPPMOD MODULE(MYLIB/IMQWRDL) +
SRCSTMF('/QIBM/ProdData/mqm/samp/imqwrld.cpp') +
INCDIR('/QIBM/ProdData/mqm/inc') DFTCHAR(*SIGNED) +
TERASPACE(*YES)
```

The source for the C++ sample programs can be found in /QIBM/ProdData/mqm/samp and the include files in /QIBM/ProdData/mqm/inc.

Alternatively, the source can be found in library SRCFILE(QCPPSRC/LIB) SRCMBR(IMQWRDL).

- b. Bind this with IBM MQ-supplied service programs to produce a program object:

```
CRTPGM PGM(MYLIB/IMQWRDL) MODULE(MYLIB/IMQWRDL) +
BNDSRVPGM(QMQM/IMQB23I4 QMQM/IMQS23I4)
```

To build a threaded application use the re-entrant service programs:

```
CRTPGM PGM(MYLIB/IMQWRDL) MODULE(MYLIB/IMQWRDL) +
BNDSRVPGM(QMQM/IMQB23I4[_R] QMQM/IMQS23I4[_R])
```

- c. Execute the HELLO WORLD sample program, using SYSTEM.DEFAULT.LOCAL.QUEUE:

```
CALL PGM(MYLIB/IMQWRDL)
```

## Building C++ programs on Linux

Build IBM MQ C++ programs on Linux using the GNU g++ compiler.

### System p

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

### Client: System p

#### 32-bit unthreaded application

```
g++ -m32 -o imqspu32 imqspu.cpp -fsigned-char -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limq23gl
-limq23gl -lmqic
```

#### 32-bit threaded application

```
g++ -m32 -o imqspu32_r imqspu.cpp -fsigned-char -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limq23gl_r
-limq23gl_r -lmqic_r
```

#### 64-bit unthreaded application

```
g++ -m64 -o imqspu64 imqspu.cpp -fsigned-char -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-limq23gl -limq23gl -lmqic
```



### 64-bit threaded application

```
g++ -m64 -o imqsputc_r64 imqsput.cpp -fsigned-char -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-limqc23gl_r -limqb23gl_r -lmqic_r
```

## Server: System p

### 32-bit unthreaded application

```
g++ -m32 -o imqsput_32 imqsput.cpp -fsigned-char -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqs23gl
-limqb23gl -lmqm
```

### 32-bit threaded application

```
g++ -m32 -o imqsput_r32 imqsput.cpp -fsigned-char -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqs23gl_r
-limqb23gl_r -lmqm_r
```

### 64-bit unthreaded application

```
g++ -m64 -o imqsput_64 imqsput.cpp -fsigned-char -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-limqs23gl -limqb23gl -lmqm
```

### 64-bit threaded application

```
g++ -m64 -o imqsput_r64 imqsput.cpp -fsigned-char -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-limqs23gl_r -limqb23gl_r -lmqm_r
```

## IBM Z

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

## Client: IBM Z

### 32-bit unthreaded application

```
g++ -m31 -fsigned-char -o imqsputc_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib
-limqc23gl -limqb23gl -lmqic
```

### 32-bit threaded application

```
g++ -m31 -fsigned-char -o imqsputc_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib
-limqc23gl_r -limqb23gl_r -lmqic_r
-lpthread
```

### 64-bit unthreaded application

```
g++ -m64 -fsigned-char -o imqsputc_64 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-limqc23gl -limqb23gl -lmqic
```

### 64-bit threaded application

```
g++ -m64 -fsigned-char -o imqsputc_64_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-limqc23gl_r -limqb23gl_r -lmqic_r -lpthread
```

## Server: IBM Z

### 32-bit unthreaded application

```
g++ -m31 -fsigned-char -o imqsput_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib
-limqs23gl -limqb23gl -lmqm
```

### 32-bit threaded application

```
g++ -m31 -fsigned-char -o imqsput_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib
-limqs23gl_r -limqb23gl_r -lmqm_r -lpthread
```

### 64-bit unthreaded application

```
g++ -m64 -fsigned-char -o imqsput_64 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-limqs23gl -limqb23gl -lmqm
```

### 64-bit threaded application

```
g++ -m64 -fsigned-char -o imqsput_64_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-limqs23gl_r -limqb23gl_r -lmqm_r -lpthread
```

## System x (32-bit)

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

### Client: System x (32-bit)

#### 32-bit unthreaded application

```
g++ -m32 -fsigned-char -o imqsputc_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -L MQ_INSTALLATION_PATH/lib -Wl,
-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqc23gl -limqb23gl -lmqic
```

#### 32-bit threaded application

```
g++ -m32 -fsigned-char -o imqsputc_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -L MQ_INSTALLATION_PATH/lib
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqc23gl_r -limqb23gl_r
-lmqic_r -lpthread
```

#### 64-bit unthreaded application

```
g++ -m64 -fsigned-char -o imqsputc_64 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -limqc23gl -limqb23gl
-lmqic
```

#### 64-bit threaded application

```
g++ -m64 -fsigned-char -o imqsputc_64_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -limqc23gl_r -limqb23gl_r
-lmqic_r -lpthread
```

### Server: System x (32-bit)

#### 32-bit unthreaded application

```
g++ -m32 -fsigned-char -o imqsput_32 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -L MQ_INSTALLATION_PATH/lib
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqs23gl -limqb23gl -lmqm
```

#### 32-bit threaded application

```
g++ -m32 -fsigned-char -o imqsput_32_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -L MQ_INSTALLATION_PATH/lib
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqs23gl_r -limqb23gl_r
-lmqm_r -lpthread
```

#### 64-bit unthreaded application

```
g++ -m64 -fsigned-char -o imqsput_64 imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -limqs23gl -limqb23gl -lmqm
```

#### 64-bit threaded application

```
g++ -m64 -fsigned-char -o imqsput_64_r imqsput.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -L MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -limqs23gl_r -limqb23gl_r
-lmqm_r -lpthread
```

## Building C++ programs on Solaris

Build IBM MQ C++ programs on Solaris using the Sun ONE compiler.

### SPARC

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

#### Client: SPARC

##### 32-bit application

```
CC -xarch=v8plus -mt -o imqsputc_32 imqsputc.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqc23as -limqb23as
-lmqic -lsocket -lnsl -ldl
```

##### 64-bit application

```
CC -xarch=v9 -mt -o imqsputc_64 imqsputc.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqc23as -limqb23as
-lmqic -lsocket -lnsl -ldl
```

#### Server: SPARC

##### 32-bit application

```
CC -xarch=v8plus -mt -o imqspu_32 imqspu.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqs23as -limqb23as
-lmqm -lsocket -lnsl -ldl
```

##### 64-bit application

```
CC -xarch=v9 -mt -o imqspu_64 imqspu.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqs23as -limqb23as
-lmqm -lsocket -lnsl -ldl
```

### x86-64

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

#### Client: x86-64

##### 32-bit application

```
CC -xarch=386 -mt -o imqsputc_32 imqsputc.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqc23as -limqb23as
-lmqic -lsocket -lnsl -ldl
```

##### 64-bit application

```
CC -xarch=amd64 -mt -o imqsputc_64 imqsputc.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqc23as -limqb23as
-lmqic -lsocket -lnsl -ldl
```

#### Server: x86-64

##### 32-bit application

```
CC -xarch=386 -mt -o imqspu_32 imqspu.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib -R MQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqs23as -limqb23as
-lmqm -lsocket -lnsl -ldl
```

##### 64-bit application

```
CC -xarch=amd64 -mt -o imqspu_64 imqspu.cpp -I MQ_INSTALLATION_PATH/inc
-L MQ_INSTALLATION_PATH/lib64 -R MQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqs23as -limqb23as
-lmqm -lsocket -lnsl -ldl
```

## Building C++ programs on Windows

Build IBM MQ C++ programs on Windows by using the Microsoft Visual Studio C++ compiler.

Library (.lib) files and dll files for use with 32-bit applications are installed in *MQ\_INSTALLATION\_PATH/Tools/Lib*, files for use with 64-bit applications are installed in *MQ\_INSTALLATION\_PATH/Tools/Lib64*. *MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

### Client

```
cl -MD imqspu.cpp /Feimqspu.exe imqb23vn.lib imqc23vn.lib
```

### Server

```
cl -MD imqspu.cpp /Feimqspu.exe imqb23vn.lib imqs23vn.lib
```

V 9.0.1

## C++ client libraries built by using the Microsoft Visual Studio 2015 compiler

V 9.0.1

From Version 9.0.1, IBM MQ provides C++ client libraries that are built with the Microsoft Visual Studio 2015 C++ compiler. Applications that are built by using a release of IBM MQ Version 9.0.1 or later can use these libraries. These libraries are provided in addition to the existing IBM MQ Version 9.0.1 C++ libraries that are built with the Microsoft Visual Studio 2012 C++ compiler.

Both 32-bit and 64-bit versions of the IBM MQ C++ libraries are provided. While the 32-bit libraries are installed under the *bin\vs2015* folder, the 64-bit libraries are installed under the *bin64\vs2015* folders.

By default, IBM MQ is configured to use the Microsoft Visual Studio 2012 libraries. To use the Microsoft Visual Studio 2015 libraries you must set the *MQ\_PREFIX\_VS\_LIBRARIES* environment variable by using the **setmqenv**, or **setmqinst** command.

You can compile your own messaging applications with Microsoft Visual Studio 2017 and link them against the supplied IBM MQ C/C++ Microsoft Visual Studio 2015 libraries.

## Using IBM MQ with the Microsoft Visual Studio 2015 C++ compiler.

In Microsoft Visual Studio 2015, the Microsoft merge modules installed with IBM MQ no longer contain the entire C runtime code.

To use the Microsoft Visual Studio 2015 C++ compiler, you must install the Microsoft Knowledge Base update, KB3118401, from the Windows 10 Universal C Runtime web page, if you are using a version of Windows before Windows 10.

This page includes system requirements and installation instructions.

If you do not install KB3118401, C++ programs built against Microsoft Visual Studio 2015 (both IBM and those for your own enterprise) will fail to launch, usually with the following message appearing:

```
The program can't start because api-ms-win-crt-runtime-|1-1-0.dll
is missing from your computer. Try reinstalling the program to
fix this problem.
```

## Using differently named IBM MQ C++ libraries

From Version 8.0.0, Fix Pack 4, IBM MQ provides some additional C++ client libraries that are named differently. These libraries are built with the Microsoft Visual Studio 2012 C++ compiler. Microsoft Visual Studio 2015 libraries are also available. These libraries are provided in addition to the existing C++ libraries that are also built with the Microsoft Visual Studio 2012 or Microsoft Visual Studio 2015 C++

compiler. Since these additional IBM MQ C++ libraries have different names, you can run IBM MQ C++ applications built by using IBM MQ C++ and compiled with Microsoft Visual Studio 2012 and earlier versions of the product on the same machine.

The additional Microsoft Visual Studio 2012 libraries are named as follows:

- imqb23vnvs2012.d11
- imqc23vnvs2012.d11
- imqs23vnvs2012.d11
- imqx23vnvs2012.d11

The additional Microsoft Visual Studio 2015 libraries are named as follows:

- imqb23vnvs2015.d11
- imqc23vnvs2015.d11
- imqs23vnvs2015.d11
- imqx23vnvs2015.d11

Both 32-bit and 64-bit versions of these libraries are provided. The 32-bit libraries are installed under the bin folder, and the 64-bit libraries are installed under the bin64 folder. Corresponding import libraries are installed under the Tools\lib and Tools\lib64 directories.

If your application uses imq\*vs2012.lib files, you must compile it using the Microsoft Visual Studio 2012 compiler. To run IBM MQ C++ applications that are compiled with Microsoft Visual Studio 2012 and applications that are compiled with an earlier version of the product on the same machine, the PATH environment variable must be prefixed as shown in the following examples:

- For 32-bit applications:  
SET PATH=*installation folder*\bin\vs2008;%PATH%
- For 64-bit applications:  
SET PATH=*installation folder*\bin64\vs2008;%PATH%

#### **Related information:**

Changes from IBM MQ Version 8.0 on Windows

## **Building C++ programs on z/OS Batch, RRS Batch and CICS**

Build IBM MQ C++ programs on z/OS for the Batch, RRS batch or CICS environments and run the sample programs.

You can write C++ programs for three of the environments that IBM MQ for z/OS supports:

- Batch
- RRS batch
- CICS

## **Compile, prelink and link**

Create an z/OS application by compiling, pre-linking, and link-editing your C++ source code.

IBM MQ C++ for z/OS is implemented as z/OS DLLs for the IBM C++ for z/OS language. Using DLLs, you concatenate the supplied definition sidedecks with the compiler output at pre-link time. This allows the linker to check your calls to the IBM MQ C++ member functions.

**Note:** There are three sets of sidedecks for each of the three environments.

To build an IBM MQ for z/OS C++ application, create and run JCL. Use the following procedure:

1. If your application runs under CICS, use the CICS-supplied procedure to translate CICS commands in your program.

In addition, for CICS applications you need to:

- a. Add the SCSQLOAD library to the DFHRPL concatenation.
  - b. Define the CSQCAT1 CEDA group using the member IMQ4B100 in the SCSQPROC library.
  - c. Install CSQCAT1.
2. Compile the program to produce object code. The JCL for your compilation must include statements that make the product data definition files available to the compiler. The data definitions are supplied in the following IBM MQ for z/OS libraries:
- **thlqual.SCSQC370**
  - **thlqual.SCSQHPPS**
- Be sure to specify the /cxx compiler option.

**Note:** The name **thlqual** is the high level qualifier of the IBM MQ installation library on z/OS.

3. Pre-link the object code created in step 2, including the following definition sidedecks, which are supplied in **thlqual.SCSQDEFS**:
- a. imqs23dm and imqb23dm for batch
  - b. imqs23dr and imqb23dr for RRS batch
  - c. imqs23dc and imqb23dc for CICS
- These are the corresponding DLLs.
- a. imqs23im and imqb23im for batch
  - b. imqs23ir and imqb23ir for RRS batch
  - c. imqs23ic and imqb23ic for CICS
4. Link-edit the object code created in step 3, to produce a load module, and store it in your application load library.

To run batch or RRS batch programs, include the libraries **thlqual.SCSQAUTH** and **thlqual.SCSQLOAD** in the STEPLIB or JOBLIB data set concatenation.

To run a CICS program, first get your system administrator to define it to CICS as an IBM MQ program and transaction. You can then run it in the usual way.

## Run the sample programs

The programs are described in “C++ sample programs” on page 798.

The sample applications are supplied in source form only. The files are:

*Table 106. z/OS sample program files*

Sample	Source program (in library <b>thlqual.SCSQCPPS</b> )	JCL (in library <b>thlqual.SCSQPROC</b> )
HELLO WORLD	imqwrlr	imqwrlr
SPUT	imqsput	imqsputr
SGET	imqsget	imqsgetr

To run the samples, compile and link-edit them as with any C++ program (see “Building C++ programs on z/OS Batch, RRS Batch and CICS” on page 821 ). Use the supplied JCL to construct and run a batch job. You must initially customize the JCL, by following the commentary included with it.

## Building C++ programs on z/OS UNIX System Services

Build IBM MQ C++ programs on z/OS for Unix System Services.

To build an application under the UNIX System Services shell, you must give the compiler access to the IBM MQ include files (located in `thlqual.SCSQC370` and `hlqual.SCSQHPPS`), and link against two of the DLL sidedecks (located in `thlqual.SCSQDEFS`). At runtime, the application needs access to the IBM MQ data sets `thlqual.SCSQLOAD`, `thlqual.SCSQAUTH`, and one of the language specific data sets, such as `thlqual.SCSQANLE`<sup>3</sup>.

### Compiling

1. Copy the sample into the HFS using the TSO **oput** command, or use FTP. The rest of this example assumes that you have copied the sample into a directory called `/u/fred/sample`, and named it `imqwrld.cpp`.
2. Log into the UNIX System Services shell, and change to the directory where you placed the sample.
3. Set up the C++ compiler so that it can accept the DLL sidedeck and `.cpp` files as input:

```
/u/fred/sample:> export _CXX_EXTRA_ARGS=1
/u/fred/sample:> export _CXX_CXXSUFFIX="cpp"
```
4. Compile and link the sample program. The following command links the program with the batch sidedecks; the RRS batch sidedecks can be used instead. The `\` character is used to split the command over more than one line. Do not enter this character; enter the command as a single line:

```
/u/fred/sample:> c++ -o imqwrld -I "'thlqual.SCSQC370'" \
-I "'thlqual.SCSQHPPS'" imqwrld.cpp \
"'thlqual.SCSQDEFS(IMQS23DM)'" "'thlqual.SCSQDEFS(IMQB23DM)'"
```

For more information on the TSO **oput** command, refer to the *z/OS UNIX System Services Command Reference*.

You can also use the `make` utility to simplify building C++ programs. Here is a sample makefile to build the HELLO WORLD C++ sample program. It separates the compiling and linking stages. Set up the environment as in step 3 before running `make`.

```
flags = -I "'thlqual.SCSQC370'" -I "'thlqual.SCSQHPPS'"
decks = "'thlqual.SCSQDEFS(IMQS23DM)'" "'thlqual.SCSQDEFS(IMQB23DM)'"

imqwrld: imqwrld.o
 c++ -o imqwrld imqwrld.o $(decks)

imqwrld.o: imqwrld.cpp
 c++ -c -o imqwrld $(flags) imqwrld.cpp
```

Refer to *z/OS UNIX System Services Programming Tools* for more information on using `make`.

### Running

1. Log into the UNIX System Services shell, and change to the directory where you built the sample.
2. Set up the `STEPLIB` environment variable to include the IBM MQ data sets:

```
/u/fred/sample:> export STEPLIB=$STEPLIB:thlqual.SCSQLOAD
/u/fred/sample:> export STEPLIB=$STEPLIB:thlqual.SCSQAUTH
/u/fred/sample:> export STEPLIB=$STEPLIB:thlqual.SCSQANLE
```
3. Run the sample:

```
/u/fred/sample:> ./imqwrld
```

---

3. You can link with any of the sidedecks listed in "Pre-link the object code to run your UNIX system service in any of the three environments, "Building C++ programs on z/OS Batch, RRS Batch and CICS" on page 821

---

## Developing JMS and Java applications

IBM MQ provides two Java language interfaces: IBM MQ classes for Java Message Service and IBM MQ classes for Java.

Within IBM MQ there are two alternative APIs for use in Java applications:

### IBM MQ classes for JMS

IBM MQ classes for Java Message Service (JMS) is the JMS provider that is supplied with IBM MQ. The Java Platform, Enterprise Edition Connector Architecture (JCA) provides a standard way of connecting applications running in a Java EE environment to an Enterprise Information System (EIS) such as IBM MQ or Db2.

### IBM MQ classes for Java

IBM MQ classes for Java enable you to use IBM MQ in a Java environment. IBM MQ classes for Java allow a Java application to connect to IBM MQ as an IBM MQ client, or connect directly to an IBM MQ queue manager.

**Note:** IBM MQ classes for Java are functionally stabilized at the level shipped in IBM MQ Version 8.0. for more information, see *Stabilization of IBM MQ classes for Java*.

## Using IBM MQ classes for JMS

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the `javax.jms` package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

The JMS specification defines a set of interfaces that applications can use to perform messaging operations. The latest version of the specification is JMS 2.0. The `javax.jms` package defines the JMS interfaces, and a JMS provider implements these interfaces for a specific messaging product. IBM MQ classes for JMS is a JMS provider that implements the JMS interfaces for IBM MQ.

The JMS specification expects `ConnectionFactory` and `Destination` objects to be administered objects. An administrator creates and maintains administered objects in a central repository, and a JMS application retrieves these objects using the Java Naming Directory Interface (JNDI). IBM MQ classes for JMS supports the use of administered objects, and an administrator can use either the IBM MQ JMS administration tool or IBM MQ Explorer to create and maintain administered objects.

IBM MQ classes for JMS also provides two sets of extensions to the JMS API. The main focus of these extensions concerns creating and configuring connection factories and destinations dynamically at run time, but the extensions also provide function that is not directly related to messaging, such as function for problem determination.

### The IBM MQ JMS extensions

Previous releases of IBM MQ classes for JMS contain extensions that are implemented in objects such as `MQConnectionFactory`, `MQQueue`, and `MQTopic` objects. These objects have properties and methods that are specific to IBM MQ. The objects can be administered objects, or an application can create the objects dynamically at run time. This release of IBM MQ classes for JMS maintains these extensions, which are now known as the IBM MQ JMS extensions. You can continue to use, without change, any applications that use these extensions.

### The IBM JMS extensions

This release of IBM MQ classes for JMS provides a more generic set of extensions to the JMS API, which are not specific to IBM MQ as the messaging system. These extensions are known as the IBM JMS extensions and have the following broad objectives:

- To provide a greater level of consistency across IBM JMS providers
- To make it easier to write a bridge application between two IBM messaging systems
- To make it easier to port an application from one IBM JMS provider to another



The extensions provide function that is similar to that provided in IBM Message Service Client for C/C++ and IBM Message Service Client for .NET.

From IBM MQ Version 8.0, the IBM MQ classes for JMS are built with Java 7.

The Java 7 runtime environment supports running earlier class file versions.

**Related concepts:**

“The JMS model” on page 871

The JMS model defines a set of interfaces that Java applications can use to perform messaging operations. IBM MQ classes for JMS, as a JMS provider, defines how JMS objects are related to IBM MQ concepts. The JMS specification expects certain JMS objects to be administered objects. JMS 2.0 introduces a simplified API, while also retaining the classic API, from JMS 1.1.

“Using JMS 2.0 functionality” on page 996

JMS 2.0 introduces several new areas of functionality to the IBM MQ classes for JMS.

**Related information:**

IBM MQ Java language interfaces

## **Why should I use IBM MQ classes for JMS?**

Using IBM MQ classes for JMS has a number of advantages including being able to reuse any existing JMS skills in your organization, and applications being more independent from the JMS provider and the underlying IBM MQ configuration.

IBM MQ classes for JMS is one of two alternative APIs that Java applications can for use to access IBM MQ resources. The other API is IBM MQ classes for Java. Although existing applications that use the IBM MQ classes for Java continue to be fully supported, new applications should use the IBM MQ classes for JMS (see “Choice of API” on page 826).

## **Summary of advantages of using IBM MQ classes for JMS**

Using IBM MQ classes for JMS allows you to reuse existing JMS skills and provide application independence.

- You can reuse JMS skills.

IBM MQ classes for JMS is a JMS provider that implements the JMS interfaces for IBM MQ as the messaging system. If your organization is new to IBM MQ, but already has JMS application development skills, you might find it easier to use the familiar JMS API to access IBM MQ resources rather than one of the other APIs provided with IBM MQ.

- JMS is an integral part of Java Platform, Enterprise Edition (Java EE).

JMS is the natural API to use for messaging on the Java EE platform. Every application server that is Java EE compliant must include a JMS provider. You can use JMS in application clients, servlets, Java Server Pages (JSPs), enterprise Java beans (EJBs), and message driven beans (MDBs). Note in particular that Java EE applications use MDBs to process messages asynchronously, and all messages are delivered to MDBs as JMS messages.

- Connection factories and destinations can be stored as JMS administered objects in a central repository rather than being hard-coded into an application.

An administrator can create and maintain JMS administered objects in a central repository, and IBM MQ classes for JMS applications can retrieve these objects by using the Java Naming Directory Interface (JNDI). JMS connection factories and destinations encapsulate IBM MQ-specific information such as queue manager names, channel names, connection options, queue names, and topic names. If connection factories and destinations are stored as administered objects, this information is not hard-coded into an application. This arrangement therefore provides the application with a degree of independence from the underlying IBM MQ configuration.

- JMS is an industry standard API that can provide application portability.

A JMS application can use JNDI to retrieve connection factories and destinations that are stored as administered objects, and use only the interfaces that are defined in the `javax.jms` package to perform messaging operations. The application is then entirely independent of any JMS provider, such as IBM MQ classes for JMS, and can be ported from one JMS provider to another without any change to the application. If JNDI is not available in a particular application environment, an IBM MQ classes for JMS application can use extensions to the JMS API to create and configure connection factories and destinations dynamically at run time. The application is then completely self-contained, but is tied to IBM MQ classes for JMS as the JMS provider.

- Bridge applications might be easier to write by using JMS.

A bridge application is an application that receives messages from one messaging system and sends them to another messaging system. Writing a bridge application can be complicated by using product-specific APIs and message formats. Instead, you can write a bridge application by using two JMS providers, one for each messaging system. The application then uses only one API, the JMS API, and processes only JMS messages.

## Deployable environments

To provide integration with a Java EE application server, the Java EE standards require messaging providers to supply a resource adapter. Following the Java EE Connector Architecture (JCA) specification, IBM MQ provides a resource adapter that uses JMS to provide messaging functions within any certified Java EE environment.

While it has been possible to use the IBM MQ classes for Java inside Java EE, this API is not engineered or optimized for this purpose. See the IBM technote [Using WebSphere MQ Java Interfaces in J2EE/JEE Environments](#) for details of IBM MQ classes for Java considerations within Java EE.

Outside of the Java EE environment, OSGi and JAR files are provided, making it easier for you to obtain just the IBM MQ classes for JMS. These JAR files are now more readily deployable either stand-alone or within software management frameworks such as Maven. For more information see the IBM technote [Obtaining the WebSphere MQ classes for JMS](#).

## Choice of API






New applications should use the IBM MQ classes for JMS rather than IBM MQ classes for Java.

IBM MQ classes for JMS provide access to both the point-to-point and publish/subscribe messaging features of IBM MQ. As well as sending JMS messages that provide support for the JMS standard messaging model, applications can also send and receive messages without additional headers and so can inter-operate with other IBM MQ applications, for example, C MQI applications. Full control of the MQMD and MQ message payloads are available. Further IBM MQ features such as message streaming, asynchronous put and report messages are also available. Using the supplied PCF helper classes, IBM MQ PCF administration messages can be sent and received through the JMS API and can be used to administer queue managers.

Features that have recently been added to IBM MQ, such as asynchronous consume and automatic reconnection, are not available in the IBM MQ classes for Java, but are available in the IBM MQ classes for JMS. Existing applications that use the IBM MQ classes for Java continue to be fully supported.

If you need access to IBM MQ features that are not available through the IBM MQ classes for JMS, you can raise a Request for Enhancement (RFE). IBM can then advise whether the implementation is possible in the IBM MQ classes for JMS implementation, or whether there is a best practice that can be followed. For additional messaging features, as IBM is a contributor to the open standard, these features can be raised as part of the JCP process.

## Related information:

-  IBM RFE Submission Process
  -  JMS Java Specification Review Process
  -  Using WebSphere MQ Java Interfaces in J2EE/JEE Environments
  -  Obtaining the WebSphere MQ classes for JMS
  -  Using JMS to send PCF messages
- Tracing IBM MQ classes for JMS applications  
Java and JMS troubleshooting



## Prerequisites for IBM MQ classes for JMS


This topic tells you what you need to know before using IBM MQ classes for JMS. To develop and run IBM MQ classes for JMS applications, you need certain software components as prerequisites.


For information about the prerequisites for IBM MQ classes for JMS, see System Requirements for IBM MQ.

To develop IBM MQ classes for JMS applications, you need a Java SE Software Development Kit (SDK). For details of the JDKs supported by your operating system, see System Requirements for IBM MQ.

To run IBM MQ classes for JMS applications, you need the following software components:

- An IBM MQ queue manager.
- A Java runtime environment (JRE), for each system on which you run applications.
-  For IBM i, Qshell, which is option 30 of the operating system.
-  For z/OS, UNIX and Linux System Services (USS).

 The IBM JSSE provider includes a FIPS certified cryptographic provider, so can be programmatically configured for FIPS 140-2 compliance ready for immediate use. Therefore, FIPS 140-2 compliance can be supported directly from IBM MQ classes for Java and IBM MQ classes for JMS.

 Oracle's JSSE provider can have a FIPS certified cryptographic provider that is configured into it (for example, see the discussion here: FIPS 140 Compliant Mode for SunJSSE), but this is not ready for immediate use and is not available for programmatic configuration. Therefore, in this case, IBM MQ classes for Java and IBM MQ classes for JMS cannot enable FIPS 140-2 compliance directly. You might be able to manually enable such compliance (again see the discussion at FIPS 140 Compliant Mode for SunJSSE for some pointers) but IBM cannot currently provide guidance on this.

You can use Internet Protocol Version 6 (IPv6) addresses in your IBM MQ classes for JMS applications if IPv6 addresses are supported by your Java virtual machine (JVM) and the TCP/IP implementation on your operating system. The IBM MQ JMS administration tool (see Configuring JMS objects using the administration tool ) also accepts IPv6 addresses.

The IBM MQ JMS administration tool and IBM MQ Explorer use the Java Naming Directory Interface (JNDI) to access a directory service, which stores administered objects. IBM MQ classes for JMS applications can also use JNDI to retrieve administered objects from a directory service. A service provider is code that provides access to a directory service by mapping JNDI calls to the directory service. A file system service provider in the files `fscontext.jar` and `providerutil.jar` is supplied with IBM MQ classes for JMS. The file system service provider provides access to a directory service based on the local file system.

If you intend to use a directory service based on an LDAP server, you must install and configure an LDAP server, or have access to an existing LDAP server. In particular, you must configure the LDAP server to store Java objects. For information about how to install and configure your LDAP server, see the documentation that is supplied with the server.

## Preparing JMS programs for the IBM MQ client for HP Integrity NonStop Server:

This topic explains what you need to know before you develop and run JMS programs for the IBM MQ client for HP Integrity NonStop Server.

The IBM MQ classes for JMS are installed as part of the IBM MQ client for HP Integrity NonStop Server installation. For details of a summary of the contents of the installation, see File system.

Some aspects of client functionality are specific to the host operating system. For more information about the supported features for the IBM MQ client for HP Integrity NonStop Server, see IBM MQ client for HP Integrity NonStop Server supported environments and features.

### Prerequisites

To build and run JMS applications, the *HP Integrity NonStop Server for Java* component must be installed and available.

### Setup

For information about setting up the environment to run and build applications in which you can use the IBM MQ classes for JMS, see “Setting environment variables” on page 833.

For information about the steps required to configure a queue manager to accept connections from client applications, see “Post installation setup for IBM MQ classes for JMS applications” on page 850.

For information about validating your IBM MQ classes for JMS environment, see “The point-to-point IVT for IBM MQ classes for JMS” on page 853.

### Writing applications

For more information about writing JMS applications, see “Writing IBM MQ classes for JMS applications” on page 871.

For more information about using the IBM MQ JMS administration tool, see Configuring JMS objects using the administration tool.

### Samples

Sample applications are provided in the following subdirectory of the installation: `opt/mqm/samp/jms`.

For more information about the configuration steps required to run the samples, see “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470.

### Problem solving

For information about solving problems, see “Solving problems with IBM MQ classes for JMS” on page 863.

## Installation and configuration of IBM MQ classes for JMS

This section describes the directories and files that are created when you install IBM MQ classes for JMS and tells you how to configure IBM MQ classes for JMS after installation.

### Related concepts:

“What is installed for IBM MQ classes for JMS” on page 830

A number of files and directories are created when you install IBM MQ classes for JMS. On Windows, some configuration is performed during installation by automatically setting environment variables. On other platforms, and in certain Windows environments, you must set environment variables before you can run IBM MQ classes for JMS applications.

“Running IBM MQ classes for JMS applications under the Java Security Manager” on page 845

IBM MQ classes for JMS can run with the Java security manager enabled. To run applications successfully with the Java Security Manager enabled, you must configure your Java virtual machine (JVM) with a suitable policy configuration file.

“Using the IBM MQ resource adapter” on page 1092

The resource adapter allows applications that are running in an application server to access IBM MQ resources. It supports inbound and outbound communication.

“Post installation setup for IBM MQ classes for JMS applications” on page 850

This topic tells you what authorities IBM MQ classes for JMS applications need in order to access the resources of a queue manager. It also introduces connection modes and describes how to configure a queue manager so that applications can connect in client mode.

“The point-to-point IVT for IBM MQ classes for JMS” on page 853

A point-to-point installation verification test (IVT) program is supplied with IBM MQ classes for JMS. The program connects to a queue manager in either bindings or client mode, sends a message to the queue called SYSTEM.DEFAULT.LOCAL.QUEUE, and then receives the message from the queue. The program can create and configure all the objects that it requires dynamically at run time, or it can use JNDI to retrieve administered objects from a directory service.

“The publish/subscribe IVT for IBM MQ classes for JMS” on page 857

A publish/subscribe installation verification test (IVT) program is supplied with IBM MQ classes for JMS. The program connects to a queue manager in either bindings or client mode, subscribes to a topic, publishes a message on the topic, and then receives the message that it has just published. The program can create and configure all the objects that it requires dynamically at run time, or it can use JNDI to retrieve administered objects from a directory service.

“Configuring the resource adapter for outbound communication” on page 1121

To configure outbound communication, define the properties of a ConnectionFactory object and an administered destination object.

“Support for OSGi” on page 862

OSGi provides a framework that supports the deployment of applications as bundles. Nine OSGi bundles are supplied as part of the IBM MQ classes for JMS.

“Solving problems with IBM MQ classes for JMS” on page 863

You can investigate problems by running the installation verification programs, and by using the trace and log facilities.

### Related tasks:

“Verifying the resource adapter installation” on page 1135

The installation verification test (IVT) program for the IBM MQ resource adapter is supplied as an EAR file. To use the program, you must deploy it and define some objects as JCA resources.

### Related reference:

“Scripts provided with IBM MQ classes for JMS” on page 861

A number of scripts are provided to assist with common tasks that need to be performed when using IBM MQ classes for JMS.

### Related information:

Problem determination for the IBM MQ resource adapter

## What is installed for IBM MQ classes for JMS:

A number of files and directories are created when you install IBM MQ classes for JMS. On Windows, some configuration is performed during installation by automatically setting environment variables. On other platforms, and in certain Windows environments, you must set environment variables before you can run IBM MQ classes for JMS applications.

For most operating systems, the IBM MQ classes for JMS are installed as an optional component when you install IBM MQ.

**NSS Client** For the IBM MQ client for HP Integrity NonStop Server, the IBM MQ classes for JMS are installed by default.

For more information about installing IBM MQ, see:

**Multi** Installing IBM MQ

**z/OS** Installing IBM MQ for z/OS

### Important:

- Apart from the relocatable JAR files described in this topic, copying the IBM MQ classes for JMS JAR files or native libraries to other machines, or to a different location on a machine where the IBM MQ classes for JMS have been installed, is not supported.
- In addition, including the `com.ibm.mq.allclient.jar` file, or the IBM MQ classes for JMS, within application archives (such as enterprise application archives, or EAR files), is not supported.

You should therefore avoid bundling IBM MQ jar files in your applications (EAR files on WebSphere Application Server), otherwise you might encounter unexpected issues associated with running back-level, unpatched code.

### Installation directories

Table 107 shows where the IBM MQ classes for JMS files are installed on each platform.

Table 107. IBM MQ classes for JMS installation directories

Platform	Directory
<b>AIX</b> UNIX and Linux	<code>MQ_INSTALLATION_PATH/java</code>
<b>Windows</b> Windows	<code>MQ_INSTALLATION_PATH\java</code>
<b>IBM i</b> IBM i	<code>/QIBM/ProdData/mqm/java</code>
<b>z/OS</b> z/OS	<code>MQ_INSTALLATION_PATH/mqm/V8R0M0/java</code>
<b>NSS Client</b> IBM MQ client for HP Integrity NonStop Server	<code>MQ_INSTALLATION_PATH/opt/mqm/java</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

The installation directory includes the following content:

- The IBM MQ classes for JMS JAR files, which are in the `MQ_INSTALLATION_PATH\java\lib` directory.
- The IBM MQ native libraries, which are used by applications that use the Java Native Interface. The 32-bit native libraries are installed into the `MQ_INSTALLATION_PATH\java\lib` directory and the 64-bit native libraries can be found in the `MQ_INSTALLATION_PATH\java\lib64` directory.

For more information about the IBM MQ native libraries, see “Configuring the Java Native Interface (JNI) libraries” on page 835.

- Additional scripts that are described in “Scripts provided with IBM MQ classes for JMS” on page 861. These scripts are in the `MQ_INSTALLATION_PATH\java\bin` directory.
- The specifications of the IBM MQ classes for JMS API. The Javadoc tool has been used to generate the HTML pages that contain the specifications of the API.

The HTML pages are in the `MQ_INSTALLATION_PATH\java\doc\WMQJMSClasses` directory:

- **ULW** On UNIX, Linux, and Windows, this subdirectory contains the individual HTML pages.
- **IBM i** On IBM i, the HTML pages are in a file called `wmqjms_javadoc.jar`.
- **z/OS** On z/OS, the HTML pages are in a file called `wmqjms_javadoc.jar`.
- Support for OSGi. OSGi bundles are installed in the `java\lib\OSGi` directory and described in “Support for OSGi” on page 862.
- The IBM MQ resource adapter, which can be deployed into any Java Platform, Enterprise Edition 7 (Java EE 7) compliant application server.  
The IBM MQ resource adapter is in the `MQ_INSTALLATION_PATH\java\lib\jca` directory; for more information, see “Using the IBM MQ resource adapter” on page 1092.
- **Windows** On Windows, symbols that can be used for debugging are installed in the `MQ_INSTALLATION_PATH\java\lib\symbols` directory.

The installation directory also includes some files that belong to other IBM MQ components:

- The IBM MQ transport for SOAP, which provides a JMS transport for SOAP, is installed into the `MQ_INSTALLATION_PATH\java\lib\soap` directory. For further information on IBM MQ transport for SOAP, see “Developing web services with IBM MQ transport for SOAP” on page 1218.  
From Version 9.0, the IBM MQ transport for SOAP is deprecated.

**V 9.0.0.3** **V 9.0.5** The `JSON4J.jar` file and `com.ibm.msg.client.mqlight` package are not needed by the IBM MQ classes for Java and IBM MQ classes for JMS. From IBM MQ Version 9.0.0, Fix Pack 3 and IBM MQ Version 9.0.5, the following changes are therefore made to the `com.ibm.mq.allclient.jar` file :

- The reference to `JSON4J.jar` file is removed from the class path statement within the manifest file for the `com.ibm.mq.allclient.jar` file.
- The package `com.ibm.msg.client.mqlight` is no longer included inside the `com.ibm.mq.allclient.jar` file.

## Sample applications

Some sample applications are supplied with IBM MQ classes for JMS. Table 108 shows where the sample applications are installed on each platform.

Table 108. Samples directories

Platform	Directory
<b>AIX</b> UNIX and Linux	<code>MQ_INSTALLATION_PATH/samp/jms</code>
<b>Windows</b> Windows	<code>MQ_INSTALLATION_PATH\tools\jms</code>
<b>IBM i</b> IBM i	<code>/QIBM/ProdData/mqm/java/samples/jms</code>
<b>z/OS</b> z/OS	<code>MQ_INSTALLATION_PATH/mqm/V8R0M0/java/samples/jms</code>

Table 108. Samples directories (continued)

Platform	Directory
<b>NSS Client</b> IBM MQ client for HP Integrity NonStop Server	<code>MQ_INSTALLATION_PATH/opt/mqm/samp/jms</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

After installation, you might need to perform some configuration tasks to compile and run applications.

“Setting environment variables” on page 833 describes the classpath that is required to run simple IBM MQ classes for JMS applications. This topic also describes additional JAR files that need to be referenced in special circumstances and the environment variables that you must set to run the scripts that are provided with IBM MQ classes for JMS.

To control properties, such as tracing and logging of an application, you need to provide a configuration properties file. The IBM MQ classes for JMS configuration properties file is described in “The IBM MQ classes for JMS configuration file” on page 837.

### Relocatable JAR files

Within an enterprise, the following files can be moved to systems that need to run IBM MQ classes for JMS:

- `-com.ibm.mq.allclient.jar`
- `-com.ibm.mq.traceControl.jar`
- `-jms.jar`
- `-fscontext.jar`
- `-providerutil.jar`
- The Bouncy Castle security provider and CMS support JAR files

The `fscontext.jar` and `providerutil.jar` files are required if your application performs JNDI lookups using a file system context.

The Bouncy Castle security provider and CMS support JAR files are required. For more information, see Support for non-IBM JREs.

The file `com.ibm.mq.allclient.jar` contains the IBM MQ classes for JMS, the IBM MQ classes for Java, and the PCF and Headers Classes. If you move this file to a new location, make sure that you take steps to keep this new location maintained with new IBM MQ Fix Packs. Also, make sure that the use of this file is made known to IBM Support if you are getting an interim fix.

To determine the version of the file `com.ibm.mq.allclient.jar`, use the following command:

```
java -jar com.ibm.mq.allclient.jar
```

The following example shows some sample output from this command:

```
C:\Program Files\IBM\MQ_1\java\lib>java -jar com.ibm.mq.allclient.jar
Name: Java Message Service Client
Version: 9.0.0.0
Level: p000-L140428.1
Build Type: Production
Location: file:/C:/Program Files/IBM/MQ_1/java/lib/com.ibm.mq.allclient.jar

Name: WebSphere MQ classes for Java Message Service
Version: 9.0.0.0
Level: p000-L140428.1
Build Type: Production
```



Location: file:/C:/Program Files/IBM/MQ\_1/java/lib/com.ibm.mq.allclient.jar

Name: WebSphere MQ JMS Provider  
 Version: 9.0.0.0  
 Level: p000-L140428.1 mqjbn=p000-L140428.1  
 Build Type: Production  
 Location: file:/C:/Program Files/IBM/MQ\_1/java/lib/com.ibm.mq.allclient.jar

Name: Common Services for Java Platform, Standard Edition  
 Version: 9.0.0.0  
 Level: p000-L140428.1  
 Build Type: Production  
 Location: file:/C:/Program Files/IBM/MQ\_1/java/lib/com.ibm.mq.allclient.jar

The file com.ibm.mq.traceControl.jar is used to dynamically control trace for IBM MQ classes for JMS applications. For more information, see Controlling trace in a running process by using IBM MQ classes for Java and IBM MQ classes for JMS.

**Related information:**

Problems in deploying the resource adapter

*Setting environment variables:*

Before you can compile and run IBM MQ classes for JMS applications, the setting for your CLASSPATH environment variable must include the IBM MQ classes for JMS Java archive (JAR) file. Depending on your requirements, you might need to add other JAR files to your class path. To run the scripts provided with IBM MQ classes for JMS, other environment variables must be set.

**About this task**

**Important:** Setting the Java option *-Xbootclasspath*, to include the IBM MQ classes for JMS, is not supported.

To compile and run IBM MQ classes for JMS applications, use the CLASSPATH setting for your platform as shown in Table 109. The setting includes the samples directory, so that you can compile and run the IBM MQ classes for JMS sample applications. Alternatively, you can specify the class path on the **java** command instead of using the environment variable.

*Table 109. CLASSPATH setting to compile and run IBM MQ classes for JMS applications, including the sample applications*








Platform	CLASSPATH setting
 AIX	CLASSPATH= MQ_INSTALLATION_PATH/java/lib/com.ibm.mqjms.jar: MQ_INSTALLATION_PATH/samp/jms/samples:
 HP Integrity NonStop Server	CLASSPATH= MQ_INSTALLATION_PATH/java/lib/com.ibm.mqjms.jar: MQ_INSTALLATION_PATH/samp/jms/samples:
 HP-UX,  Linux,  Solaris	CLASSPATH= MQ_INSTALLATION_PATH/java/lib/com.ibm.mqjms.jar: MQ_INSTALLATION_PATH/samp/jms/samples:
 IBM i	CLASSPATH=/QIBM/ProdData/mqm/java/lib/com.ibm.mqjms.jar: /QIBM/ProdData/mqm/java/samples/jms/samples:
 Windows	CLASSPATH= MQ_INSTALLATION_PATH\java\lib\com.ibm.mqjms.jar; MQ_INSTALLATION_PATH\tools\jms\samples;

Table 109. CLASSPATH setting to compile and run IBM MQ classes for JMS applications, including the sample applications (continued)

Platform	CLASSPATH setting
z/OS	CLASSPATH= MQ_INSTALLATION_PATH/mqm/V8R0M0/java/lib/com.ibm.mqjms.jar: MQ_INSTALLATION_PATH/mqm/V8R0M0/java/samples/jms/samples:

MQ\_INSTALLATION\_PATH represents the high-level directory in which IBM MQ is installed.

The manifest of the JAR file com.ibm.mqjms.jar contains references to most of the other JAR files required by IBM MQ classes for JMS applications, and so you do not need to add these JAR files to your class path. These JAR files include those required by applications that use the Java Naming Directory Interface (JNDI) to retrieve administered objects from a directory service and by applications that use the Java Transaction API (JTA).

However, you must include additional JAR files in your class path in the following circumstances:

- If you are using channel exit classes that implement the channel exit interfaces defined in the com.ibm.mq package, instead of those defined in the com.ibm.mq.exits package, you must add the IBM MQ classes for Java JAR file, com.ibm.mq.jar, to your class path.
- If your application uses JNDI to retrieve administered objects from a directory service, you must also add the following JAR files to your class path:
  - fscontext.jar
  - providerutil.jar
- If your application uses the JTA, you must also add jta.jar to your class path.

**Note:** These additional JAR files are required only for compiling your applications, not for running them.

The scripts provided with IBM MQ classes for JMS use the following environment variables:

#### MQ\_JAVA\_DATA\_PATH

This environment variable specifies the directory for log and trace output.

#### MQ\_JAVA\_INSTALL\_PATH

This environment variable specifies the directory where IBM MQ classes for JMS is installed.

#### MQ\_JAVA\_LIB\_PATH

This environment variable specifies the directory where the IBM MQ classes for JMS libraries are stored, as shown in Table 110 on page 835.



### Procedure

- **Windows** On Windows, after installing IBM MQ, run the command **setmqenv**. If you do not run this command first, the following error message might appear when you are issuing a **dspmqver** command:
 


```
> V 9.0.2
AMQ8351: IBM MQ Java environment has not been configured
correctly, or the IBM MQ JRE feature has not been installed.
```

**Note:** **V 9.0.2** This message is to be expected if you did not install the IBM MQ Java Runtime Environment (JRE).

- On any other platform, set the environment variables yourself:
  - **UNIX** **NSS Client** **Linux** To set the environment variables if you are using a 32-bit JVM on UNIX, HP Integrity NonStop Server, or Linux systems, you can use the script setjmsenv.

-   To set the environment variables if you are using a 64-bit JVM on a UNIX or Linux system, you can use the script `setjmsenv64`. These scripts are in the `MQ_INSTALLATION_PATH/java/bin` directory, where `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

You can use the `setjmsenv` or `setjmsenv64` script in a variety of ways: You can use it as a basis for setting the required environment variables, as shown in the table, or add them to `.profile` using a text editor. If you have a non-typical setup, edit the script contents as necessary. Alternatively, you can run the script in every session from which JMS startup scripts are to be run. If you choose this option you need to run the script in every shell window you start, during the JMS verification process by typing `./setjmsenv` or `./setjmsenv64`.

 On IBM i, you must set the environment variable `QIBM_MULTI_THREADED` to `Y`. You can then run multithreaded applications in the same way that you run single threaded applications. See *Setting up IBM MQ with Java and JMS* for more information.

*Configuring the Java Native Interface (JNI) libraries:*

IBM MQ classes for JMS applications, that either connect to a queue manager using the bindings transport, or that connect to a queue manager using the client transport and use channel exit programs written in languages other than Java, need to be run in an environment that allows access to the Java Native Interface (JNI) libraries.

**About this task**

To set up this environment, you must configure the environment's library path so that the Java virtual machine (JVM) can load the `mqjbnd` library before you start the IBM MQ classes for JMS application.

IBM MQ provides two Java Native Interface (JNI) libraries:

**mqjbnd**








This library is used by applications that connect to a queue manager using the bindings transport. It provides the interface between the IBM MQ classes for JMS and the queue manager. The `mqjbnd` library installed with IBM MQ Version 9.0 can be used to connect to any IBM MQ Version 9.0 (or earlier) queue manager.

**mqjexitstub02**


The `mqjexitstub02` library is loaded by the IBM MQ classes for JMS when an application connects to a queue manager using the client transport and uses a channel exit program written in a language other than Java.

On certain platforms, IBM MQ installs 32-bit and 64-bit versions of these JNI libraries. The location of the libraries for each platform is shown in Table 1.

*Table 110. The location of the IBM MQ classes for JMS libraries for each platform*

Platform	Directory containing the IBM MQ classes for JMS libraries
 AIX  HP-UX  Linux  Linux (POWER, x86-64 and zSeries s390x platforms)  Solaris (x86-64 and SPARC platforms)	<code>MQ_INSTALLATION_PATH/java/lib</code> (32-bit libraries) <code>MQ_INSTALLATION_PATH/java/lib64</code> (64-bit libraries)
 Windows	<code>MQ_INSTALLATION_PATH\java\lib</code> (32-bit libraries) <code>MQ_INSTALLATION_PATH\java\lib64</code> (64-bit libraries)
 z/OS	<code>MQ_INSTALLATION_PATH/mqm/V8R0M0/java/lib</code> (31-bit and 64-bit libraries)

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.


**Note:**  On z/OS, you can use either a 31-bit or 64-bit Java virtual machine (JVM). You do not have to specify which JNI libraries to use; IBM MQ classes for JMS can determine for itself which JNI libraries to load.

## Procedure

1. Configure the JVM's `java.library.path` property, which can be done in two ways:

- By specifying the JVM argument as shown in the following example:

```
-Djava.library.path=path_to_library_directory
```

 For example, for a 64-bit JVM on Linux for a default location installation, specify:

```
-Djava.library.path=/opt/mqm/java/lib64
```

- By configuring the shell's environment such that the JVM will set up its own `java.library.path`. This path varies by platform and by the location in which you installed IBM MQ. For example, for a 64-bit JVM and a default IBM MQ installation location, you can use the following settings:



```
export LIBPATH=/usr/mqm/java/lib64:$LIBPATH
```

```
export LD_LIBRARY_PATH=/opt/mqm/java/lib64:$LD_LIBRARY_PATH
```



```
set PATH=C:\Program Files\IBM\MQ\java\lib64;%PATH%
```

An example of the exception stack that you see when the environment has not been configured correctly is as follows:

```
Caused by: com.ibm.mq.jmqi.local.LocalMQ$4: CC=2;RC=2495;
AMQ8598: Failed to load the WebSphere MQ native JNI library: 'mqjbnf'.
 at com.ibm.mq.jmqi.local.LocalMQ.loadLib(LocalMQ.java:1268)
 at com.ibm.mq.jmqi.local.LocalMQ$1.run(LocalMQ.java:309)
 at java.security.AccessController.doPrivileged(AccessController.java:400)
 at com.ibm.mq.jmqi.local.LocalMQ.initialise_inner(LocalMQ.java:259)
 at com.ibm.mq.jmqi.local.LocalMQ.initialise(LocalMQ.java:221)
 at com.ibm.mq.jmqi.local.LocalMQ.<init>(LocalMQ.java:1350)
 at com.ibm.mq.jmqi.local.LocalServer.<init>(LocalServer.java:230)
 at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method)
 at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:86)
 at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:58)
 at java.lang.reflect.Constructor.newInstance(Constructor.java:542)
 at com.ibm.mq.jmqi.JmqiEnvironment.getInstance(JmqiEnvironment.java:706)
 at com.ibm.mq.jmqi.JmqiEnvironment.getMQI(JmqiEnvironment.java:640)
 at com.ibm.msg.client.wmq.factories.WMQConnectionFactory.createV7ProviderConnection(WMQConnectionFactory.java:8437)
 ... 7 more
Caused by: java.lang.UnsatisfiedLinkError: mqjbnf (Not found in java.library.path)
 at java.lang.ClassLoader.loadLibraryWithPath(ClassLoader.java:1235)
 at java.lang.ClassLoader.loadLibraryWithClassLoader(ClassLoader.java:1205)
 at java.lang.System.loadLibrary(System.java:534)
 at com.ibm.mq.jmqi.local.LocalMQ.loadLib(LocalMQ.java:1240)
 ... 20 more
```

2. After either the 32-bit or 64-bit environment has been set up, start the IBM MQ classes for JMS application using the command:

```
java application-name
```

where *application-name* is the name of the IBM MQ classes for JMS application to be run.

An exception containing IBM MQ Reason code 2495 (MQRC\_MODULE\_NOT\_FOUND) is thrown by the IBM MQ classes for JMS if:

- The IBM MQ classes for JMS application is run in a 32-bit Java runtime environment, and a 64-bit environment has been set up for the IBM MQ classes for JMS, as the 32-bit Java runtime environment is unable to load the 64-bit Java Native Library.
- The IBM MQ classes for JMS application is run in a 64-bit Java runtime environment, and a 32-bit environment has been set up for the IBM MQ classes for JMS, as the 64-bit Java runtime environment is unable to load the 32-bit Java Native Library.

*The IBM MQ classes for JMS configuration file:*

An IBM MQ classes for JMS configuration file specifies properties that are used to configure IBM MQ classes for JMS.

**Note:** The properties defined in configuration file can also be set as JVM system properties. If a property is set both in the configuration file and as a system property, the system property takes precedence. Therefore, if required, you can override any property in the configuration file by specifying it as a system property on the **java** command.

The format of an IBM MQ classes for JMS configuration file is that of a standard Java properties file. A sample configuration file called `jms.config` is supplied in the `bin` subdirectory of the IBM MQ classes for JMS installation directory. This file documents all the supported properties and their default values.

You can choose the name and location of a IBM MQ classes for JMS configuration file. When you start your application, use a **java** command with the following format:

```
java -Dcom.ibm.msg.client.config.location= config_file_url application_name
```

In the command, *config\_file\_url* is a uniform resource locator (URL) that specifies the name and location of the IBM MQ classes for JMS configuration file. URLs of the following types are supported: `http`, `file`, `ftp`, and `jar`.

Here is an example of a **java** command:

```
java -Dcom.ibm.msg.client.config.location=file:/D:/mydir/myjms.config MyAppClass
```

This command identifies the IBM MQ classes for JMS configuration file as the file `D:\mydir\mjms.config` on the local Windows system.

When an application starts, IBM MQ classes for JMS reads the contents of the configuration file and stores the specified properties in an internal property store. If the **java** command does not identify a configuration file, or if the configuration file cannot be found, IBM MQ classes for JMS uses the default values for all the properties.

An IBM MQ classes for JMS configuration file can be used with any of the supported transports between an application and a queue manager or broker.

### **Overriding properties specified in an IBM MQ MQI client configuration file**

An IBM MQ MQI client configuration file can also specify properties that are used to configure IBM MQ classes for JMS. However, properties specified in an IBM MQ MQI client configuration file apply only when an application connects to a queue manager in client mode.

If required, you can override any attribute in a IBM MQ MQI client configuration file by specifying it as a property in a IBM MQ classes for JMS configuration file. To override an attribute in a IBM MQ MQI client configuration file, use an entry with the following format in the IBM MQ classes for JMS configuration file:

`com.ibm.mq.cfg. stanza. propName = propValue`

The variables in the entry have the following meanings:

**stanza** The name of the stanza in the IBM MQ MQI client configuration file that contains the attribute

**propName**

The name of the attribute as specified in the IBM MQ MQI client configuration file

**propValue**

The value of the property that overrides the value of the attribute specified in the IBM MQ MQI client configuration file

Alternatively, you can override an attribute in an IBM MQ MQI client configuration file by specifying the property as a system property on the **java** command. Use the preceding format to specify the property as a system property.

Only the following attributes in an IBM MQ MQI client configuration file are relevant to IBM MQ classes for JMS. If you specify or override other attributes, it has no effect. Specifically, please note that the `ChannelDefinitionFile` and `ChannelDefinitionDirectory` in the `CHANNELS` stanza of the client configuration file are not used. See “Using a client channel definition table with IBM MQ classes for JMS” on page 982 for details of how to use the CCDT with the IBM MQ classes for JMS.

Table 111. Which stanza of the client configuration file contains which attribute

Stanza	Attribute
CHANNELS stanza of the client configuration file	Put1DefaultAlwaysSync
CHANNELS stanza of the client configuration file	DefRecon
CHANNELS stanza of the client configuration file	ReconDelay
CHANNELS stanza of the client configuration file	PasswordProtection
ClientExitPath stanza of the client configuration file	ExitsDefaultPath
ClientExitPath stanza of the client configuration file	ExitsDefaultPath64
ClientExitPath stanza of the client configuration file	JavaExitsClasspath
JMQI stanza of the client configuration file	useMQCSPauthentication
MessageBuffer stanza of the client configuration file	MaximumSize
MessageBuffer stanza of the client configuration file	PurgeTime
MessageBuffer stanza of the client configuration file	UpdatePercentage
TCP stanza of the client configuration file	ClntRcvBufSize
TCP stanza of the client configuration file	ClntSndBufSize
TCP stanza of the client configuration file	Connect_Timeout
TCP stanza of the client configuration file	KeepAlive

For further details on the IBM MQ MQI client configuration see [Configuring a client using a configuration file](#)

Java Standard Environment Trace stanza:

Use the Java Standard Environment Trace Settings stanza to configure the IBM MQ classes for JMS trace facility.

**com.ibm.msg.client.commonservices.trace.outputName = *traceOutputName***

*traceOutputName* is the directory and file name to which trace output is sent.

*traceOutputName* defaults to a file named mqjms\_*PID*.trc in the current working directory where *PID* is the current process ID. If a process ID is unavailable, a random number is generated and prefixed with the letter f. To include the process ID in a file name you specify, use the string %*PID*%.

If you specify an alternative directory, it must exist, and you must have write permission for this directory. If you do not have write permission, the trace output is written to System.err.

**com.ibm.msg.client.commonservices.trace.include = *includeList***

*includeList* is a list of packages and classes that are traced, or the special values ALL or NONE.

Separate package or class names with a semicolon, ;. *includeList* defaults to ALL, and traces all packages and classes in IBM MQ classes for JMS.

**Note:** You can include a package but then exclude subpackages of that package. For example, if you include package a.b and exclude package a.b.x, the trace includes everything in a.b.y and a.b.z, but not a.b.x or a.b.x.1.

**com.ibm.msg.client.commonservices.trace.exclude = *excludeList***

*excludeList* is a list of packages and classes that are not traced, or the special values ALL or NONE.

Separate package or class names with a semicolon, ;. *excludeList* defaults to NONE, and therefore excludes no packages and classes in IBM MQ classes for JMS from being traced.

**Note:** You can exclude a package but then include subpackages of that package. For example, if you exclude package a.b and include package a.b.x, the trace includes everything in a.b.x and a.b.x.1, but not a.b.y or a.b.z.

Any package or class that is specified, at the same level, as both included and excluded is included.

**com.ibm.msg.client.commonservices.trace.maxBytes = *maxArrayBytes***

*maxArrayBytes* is the maximum number of bytes that are traced from any byte arrays.

If *maxArrayBytes* is set to a positive integer, it limits the number of bytes in a byte-array that are written out to the trace file. It truncates the byte array after writing *maxArrayBytes* out. Setting *maxArrayBytes* reduces the size of the resulting trace file, and reduces the effect of tracing on the performance of the application.

A value of 0 for this property means that none of the contents of any byte arrays are sent to the trace file.

The default value is -1, which removes any limit on the number of bytes in a byte array that are sent to the trace file.

**com.ibm.msg.client.commonservices.trace.limit = *maxTraceBytes***

*maxTraceBytes* is the maximum number of bytes that are written to a trace output file.

*maxTraceBytes* works with *traceCycles*. If the number of bytes of trace written is near to the limit, the file is closed, and a new trace output file is started.

A value of 0 means that a trace output file has zero length. The default value is -1, which means that the amount of data to be written to a trace output file is unlimited.

**com.ibm.msg.client.commonservices.trace.count = *traceCycles***

*traceCycles* is the number of trace output files to cycle through.

If the current trace output file reaches the limit specified by *maxTraceBytes*, the file is closed. Further trace output is written to the next trace output file in sequence. Each trace output file is distinguished by a numeric suffix appended to the file name. The current or most recent trace output file is `mqjms.trc.0`, the next most recent trace output file is `mqjms.trc.1`. Older trace files follow the same numbering pattern up to the limit.

The default value of *traceCycles* is 1. If *traceCycles* is 1, when the current trace output file reaches its maximum size, the file is closed and deleted. A new trace output file with the same name is started. Therefore, only one trace output file exists at a time.

**com.ibm.msg.client.commonservices.trace.parameter = *traceParameters***

*traceParameters* controls whether method parameters and return values are included in the trace.

*traceParameters* defaults to TRUE. If *traceParameters* is set to FALSE, only method signatures are traced.

**com.ibm.msg.client.commonservices.trace.startup = *startup***

There is an initialization phase of IBM MQ classes for JMS during which resources are allocated. The main trace facility is initialized during the resource allocation phase.

If *startup* is set to TRUE, startup trace is used. Trace information is produced immediately and includes the setup of all components, including the trace facility itself. Startup trace information can be used to diagnose configuration problems. Startup trace information is always written to `System.err`.

*startup* defaults to FALSE.

*startup* is checked before initialization is complete. For this reason, only specify the property on the command line as a Java system property. Do not specify it in the IBM MQ classes for JMS configuration file.

**com.ibm.msg.client.commonservices.trace.compress = *compressedTrace***

Set *compressedTrace* to TRUE to compress trace output.

The default value of *compressedTrace* is FALSE.

If *compressedTrace* is set to TRUE, trace output is compressed. The default trace output file name has the extension `.trz`. If compression is set to FALSE, the default value, the file has the extension `.trc` to indicate it is uncompressed. However if the file name for the trace output has been specified in *traceOutputName* that name is used instead; no suffix is applied to the file.

Compressed trace output is smaller than uncompressed. Because there is less I/O, it can be written out faster than uncompressed trace. Compressed tracing has less effect on the performance of IBM MQ classes for JMS than uncompressed tracing.

If *maxTraceBytes* and *traceCycles* are set, multiple compressed trace files are created in place of multiple flat files.

If IBM MQ classes for JMS ends in an uncontrolled manner, a compressed trace file might not be valid. For this reason, trace compression must only be used when IBM MQ classes for JMS closes down in a controlled manner. Only use trace compression if the problems being investigated do not cause the JVM itself to stop unexpectedly. Do not use trace compression when diagnosing problems that can result in `System.Halt()` shutdowns or abnormal, uncontrolled JVM terminations.

**com.ibm.msg.client.commonservices.trace.level = *traceLevel***

*traceLevel* specifies a filtering level for the trace. The defined trace levels are as follows:

- TRACE\_NONE: 0
- TRACE\_EXCEPTION: 1
- TRACE\_WARNING: 3
- TRACE\_INFO: 6
- TRACE\_ENTRYEXIT: 8



- TRACE\_DATA: 9
- TRACE\_ALL: Integer.MAX\_VALUE

Each trace level includes all lower levels. For example, if trace level is set at TRACE\_INFO, then any trace point with a defined level of TRACE\_EXCEPTION, TRACE\_WARNING, or TRACE\_INFO is written to the trace. All other trace points are excluded.

**com.ibm.msg.client.commonservices.trace.standalone = *standaloneTrace***

*standaloneTrace* controls whether the IBM MQ JMS client tracing service is used in a WebSphere Application Server environment.

If *standaloneTrace* is set to TRUE, the IBM MQ JMS client tracing properties are used to determine the trace configuration.

If *standaloneTrace* is set to FALSE, and the IBM MQ JMS client is running in an WebSphere Application Server container, the WebSphere Application Server trace service is used. The trace information that is generated depends upon the trace settings of the application server.

The default value of *standaloneTrace* is FALSE.

*Logging stanza:*

Use the Logging stanza to configure the IBM MQ classes for JMS log facility.

The following properties can be included in the Logging stanza:

**com.ibm.msg.client.commonservices.log.outputName = path**

The name of the log file that is used by the IBM MQ classes for JMS log facility. The default value is mqjms.log, which is written to the current working directory for the Java Runtime Environment that the IBM MQ classes for JMS are running in.

The property can take one of the following values:

- a single path name
- a comma-separated list of path names (all data is logged to all files)

Each path name can be an absolute or relative path name or:

**"stderr" or "System.err"**

Represents the standard error stream.

**"stdout" or "System.out"**

Represents the standard output stream.

**com.ibm.msg.client.commonservices.log.maxBytes**

The maximum number of bytes that are logged from any call to log message data.

**Positive integer**

Data is written up to that value of bytes per log call.

0 No data is written.

-1 Unlimited data is written (default).

**com.ibm.msg.client.commonservices.log.limit**

The maximum number of bytes that are written to any 1 log file (default is 262144).

**Positive integer**

Data is written up to that value of bytes per log file.

0 No data is written.

-1 Unlimited data is written.

**com.ibm.msg.client.commonservices.log.count**

The number of log files to cycle through. As each file reaches `com.ibm.msg.client.commonservices.trace.limit` trace will begin in the next file, the default is 3.

**Positive integer**

Number of files to cycle through.

**0** A single file.

*Java SE Specifics stanza:*

Use the Java SE Specifics stanza to configure properties that are used when the IBM MQ classes for JMS are being used in a Java Standard Edition environment.

**com.ibm.msg.client.commonservices.j2se.produceJavaCore = TRUE|FALSE**

Determines whether a JavaCore file is written immediately after the IBM MQ classes for JMS has generated an FDC file. If this is set to TRUE a JavaCore file is produced in the working directory of the Java Runtime Environment in which the IBM MQ classes for JMS are running.

**TRUE** Generate JavaCore, subject to the Java Runtime Environment's ability to do so.

**FALSE** Do not generate JavaCore; this is the default value.

*IBM MQ Properties stanza:*

Use the IBM MQ Properties stanza to set properties that affect how the IBM MQ classes for JMS interact with IBM MQ.

**com.ibm.msg.client.wmq.compat.base.internal.MQQueue.smallMsgsBufferReductionThreshold**

When an application that uses the IBM MQ classes for JMS is connecting to an IBM MQ queue manager using IBM MQ messaging provider migration mode, the IBM MQ classes for JMS uses a default buffer size of 4 KB when it receives messages. If the message that the application is trying to get is larger than 4 KB, the IBM MQ classes for JMS resizes the buffer to be large enough to accommodate the message. The larger buffer size is then used when subsequent messages are received.

This property controls when the buffer size is reduced back to 4 KB. By default, when ten consecutive messages that are less than the larger buffer size are received, the buffer size is reduced back to 4 KB. To reset the buffer size back to 4 KB every time a message is received, set the property to the value 0.

**0**

The buffer always resets to the default size.

**10**

This is the default value. The buffer will be resized after the tenth message.

**com.ibm.msg.client.wmq.receiveConversionCCSID**

When an application that is using the IBM MQ classes for JMS is connecting to an IBM MQ queue manager using IBM MQ messaging provider normal mode, the `receiveConversionCCSID` property can be set to override the default CCSID value in the MQMD structure that is used to receive messages from the queue manager. By default, the MQMD contains a CCSID field set to 1208, but this can be changed if, for example, the queue manager is unable to convert messages to this code page.

Valid values are any valid CCSID number or one of the following values:

**-1** Use the platform default.

**1208** This is the default value.

*Client-mode specifics stanza:*

Use the Client-mode specifics stanza to specify properties that are used when the IBM MQ classes for JMS connect to a queue manager that is using the CLIENT transport.

**com.ibm.mq.polling.RemoteRequestEntry**

Specifies the polling interval that the IBM MQ classes for JMS uses to check for broken connections when it is waiting for a response from a queue manager.

**Positive integer**

The number of milliseconds to wait before checking. The default value is 10000 or 10 seconds. The minimum value is 3000, and lower values are treated in the same way as this minimum value.

*Properties used to configure JMS client behavior:*

Use these properties to configure the behavior of the JMS client.

**com.ibm.mq.jms.SupportMQExtensions TRUE|FALSE**

The JMS 2.0 specification introduces changes to the way certain behaviors work. IBM MQ Version 8.0 includes the property `com.ibm.mq.jms.SupportMQExtensions`, which can be set to *TRUE*, to revert these changed behaviors back to previous implementations. Reverting the changed behaviors might be necessary some for JMS 2.0 applications, and also for some applications that use the JMS 1.1 API but run against the IBM MQ Version 8.0 IBM MQ classes for JMS.

**TRUE** The following three areas of functionality are reverted by setting `SupportMQExtensions` to *TRUE*:

**Message priority**

Messages can be assigned a priority, 0 - 9. Before JMS 2.0, messages could also use the value *-1*, indicating that a queue's default priority is used. JMS 2.0 does not allow a message priority of *-1* to be set. Turning on `SupportMQExtensions` allows the value of *-1* to be used.

**Client id**

The JMS 2.0 specification requires that non-null client ids are checked for uniqueness when they make a connection. Turning on `SupportMQExtensions`, means that this requirement is disregarded, and that a client ID can be reused.

**NoLocal**

The JMS 2.0 specification requires that when this constant is turned on, a consumer cannot receive messages that are published by the same client ID. Before JMS 2.0, this attribute was set on a subscriber to prevent it receiving messages that are published by its own connection. Turning on `SupportMQExtensions` reverts this behavior to its previous implementation.

**FALSE**

The changes of behavior are retained.

**com.ibm.msg.client.jms.ByteStreamReadOnlyAfterSend= TRUE|FALSE**

From IBM MQ Version 8.0.0, Fix Pack 2, after an application has sent a Bytes or Stream message, IBM MQ classes for JMS can set the state of the message that has just been sent to either read only, or write only.

**TRUE** The objects are set to read only after being sent. Setting this value maintains compatibility with the JMS 2.0 specification

**FALSE**

The objects are set to write only after being sent. This is the default value.

### Related concepts:

“SupportMQExtensions property” on page 1001

The JMS 2.0 specification introduces changes to the way certain behaviors work. IBM MQ Version 8.0 and later includes the property `com.ibm.mq.jms.SupportMQExtensions`, which can be set to `TRUE`, to revert these changed behaviors back to previous implementations.

*STEPLIB configuration for IBM MQ classes for JMS on z/OS:*

On z/OS, the STEPLIB used at run time must contain the IBM MQ SCSQAUTH and SCSQANLE libraries. Specify these libraries in the startup JCL or using the `.profile` file.

From UNIX and Linux System Services, you can add these using a line in your `.profile` as shown in the following code snippet, replacing `thlqual` with the high-level data set qualifier that you chose when installing IBM MQ:

```
export STEPLIB=thlqual.SCSQAUTH:thlqual.SCSQANLE:$STEPLIB
```

In other environments, you typically need to edit the startup JCL to include SCSQAUTH and SCSQANLE on the STEPLIB concatenation:

```
STEPLIB DD DSN=thlqual.SCSQAUTH,DISP=SHR
DD DSN=thlqual.SCSQANLE,DISP=SHR
```

*IBM MQ classes for JMS and software management tools:*

Software management tools such as Apache Maven can be used with the IBM MQ classes for JMS.

Many large development organizations use these tools to centrally manage repositories of third-party libraries.

The IBM MQ classes for JMS are composed of a number of JAR files. When you are developing Java language applications by using this API, an installation of either an IBM MQ Server, IBM MQ Client, or IBM MQ Client SupportPac is required on the machine where the application is being developed.

If you want to use such a tool and add the JAR files that make up the IBM MQ classes for JMS to a centrally managed repository, the following points must be observed:

- A repository or container must be made available only to developers within your organization. Any distribution outside of the organization is not permitted.
- The repository needs to contain a complete and consistent set of JAR files from a single IBM MQ release or Fix Pack.
- You are responsible for updating the repository with any maintenance provided by IBM Support.

From IBM MQ Version 8.0, the following JAR files need to be installed into the repository:

- `com.ibm.mq.allclient.jar`.
- `jms.jar` is required if you are using the IBM MQ classes for JMS.
- `fscontext.jar` is required if you are using the IBM MQ classes for JMS and accessing JMS administered objects that are stored in a file system JNDI context.
- `providerutil.jar` if you are using the IBM MQ classes for JMS and accessing JMS administered objects that are stored in a file system JNDI context.

From IBM MQ Version 9.0, the Bouncy Castle security provider and CMS support JAR files are required. For more information, see “What is installed for IBM MQ classes for JMS” on page 830 and Support for non-IBM JREs.

## Running IBM MQ classes for JMS applications under the Java Security Manager:

IBM MQ classes for JMS can run with the Java security manager enabled. To run applications successfully with the Java Security Manager enabled, you must configure your Java virtual machine (JVM) with a suitable policy configuration file.

The simplest way to create a suitable policy definition file is to change the policy configuration file supplied with your Java runtime environment (JRE). On most systems, this file is in the directory `lib/security/java.policy` relative to your JRE directory. You can edit the policy configuration file either by using your preferred editor or by using the policy tool program supplied with your JRE.

**Important:** **V 9.0.1** If you use the Java Security Manager mechanism with your application, you must grant the following permissions:

- FilePermission on any whitelist file that you use, with read permission for ENFORCEMENT mode, write permission for DISCOVER mode.
- PropertyPermission (read) on the `com.ibm.mq.jms.whitelist`, `com.ibm.mq.jms.whitelist.discover`, and `com.ibm.mq.jms.whitelist.mode` properties.

For Continuous Delivery, ClassName whitelisting is supported from IBM MQ Version 9.0.1. For more information, see “ClassName whitelisting in JMS ObjectMessage” on page 847.

**V 9.0.0.1** In the Long Term Support release, ClassName whitelisting is supported with APAR IT14385, and from IBM MQ Version 9.0.0, Fix Pack 1.

### Example policy configuration file

Here is an example of a policy configuration file that allows IBM MQ classes for JMS to run successfully under the default security manager. This file will need to be customized, to specify the locations of certain files and directories: `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed, `MQ_DATA_DIRECTORY` represents the location of the MQ data directory, and `QM_NAME` is the name of the queue manager for which access is being configured.

```
grant codeBase "file:MQ_INSTALLATION_PATH/java/lib/*" {
 //We need access to these properties, mainly for tracing
 permission java.util.PropertyPermission "user.name","read";
 permission java.util.PropertyPermission "os.name","read";
 permission java.util.PropertyPermission "user.dir","read";
 permission java.util.PropertyPermission "line.separator","read";
 permission java.util.PropertyPermission "path.separator","read";
 permission java.util.PropertyPermission "file.separator","read";
 permission java.util.PropertyPermission "com.ibm.msg.client.commonservices.log.*","read";
 permission java.util.PropertyPermission "com.ibm.msg.client.commonservices.trace.*","read";
 permission java.util.PropertyPermission "Diagnostics.Java.Errors.Destination.FileName","read";
 permission java.util.PropertyPermission "com.ibm.mq.commonservices","read";
 permission java.util.PropertyPermission "com.ibm.mq.cfg.*","read";

 //Tracing - we need the ability to control java.util.logging
 permission java.util.logging.LoggingPermission "control";
 // And access to create the trace file and read the log file - assumed to be in the current directory
 permission java.io.FilePermission "*","read,write";

 // We'd like to set up an mBean to control trace
 permission javax.management.MBeanServerPermission "createMBeanServer";
 permission javax.management.MBeanPermission "*","*";

 // We need to be able to read manifests etc from the jar files in the installation directory
 permission java.io.FilePermission "MQ_INSTALLATION_PATH/java/lib/-","read";

 //Required if mqclient.ini/mqs.ini configuration files are used
 permission java.io.FilePermission "MQ_DATA_DIRECTORY/mqclient.ini","read";
 permission java.io.FilePermission "MQ_DATA_DIRECTORY/mqs.ini","read";
}
```

```

//For the client transport type.
permission java.net.SocketPermission "*" ,"connect,resolve";

//For the bindings transport type.
permission java.lang.RuntimePermission "loadLibrary.*";

//For applications that use CCDT tables (access to the CCDT AMQCLCHL.TAB)
permission java.io.FilePermission "MQ_DATA_DIRECTORY/qmgrs/QM_NAME/@ipcc/AMQCLCHL.TAB", "read";

//For applications that use User Exits
permission java.io.FilePermission "MQ_DATA_DIRECTORY/exits/*", "read";
permission java.io.FilePermission "MQ_DATA_DIRECTORY/exits64/*", "read";
permission java.lang.RuntimePermission "createClassLoader";

//Required for the z/OS platform
permission java.util.PropertyPermission "com.ibm.vm.bitmode", "read";

// Used by the internal ConnectionFactory implementation
permission java.lang.reflect.ReflectPermission "suppressAccessChecks";

// Used for controlled class loading
permission java.lang.RuntimePermission "setContextClassLoader";

// Used to default the Application name in Client mode connections
permission java.util.PropertyPermission "sun.java.command", "read";

// Used by the IBM JSSE classes
permission java.util.PropertyPermission "com.ibm.crypto.provider.AESNITrace", "read";

//Required to determine if an IBM Java Runtime is running in FIPS mode,
//and to modify the property values status as required.
permission java.util.PropertyPermission "com.ibm.jsse2.usefipsprovider", "read,write";
permission java.util.PropertyPermission "com.ibm.jsse2.JSSEFIPS", "read,write";
//Required if an IBM FIPS provider is to be used for SSL communication.
permission java.security.SecurityPermission "insertProvider.IBMJCEFIPS";

// Required for non-IBM Java Runtimes that establish secure client
// transport mode connections using mutual TLS authentication
permission java.util.PropertyPermission "javax.net.ssl.keyStore", "read";
permission java.util.PropertyPermission "javax.net.ssl.keyStorePassword", "read";
};

```

In the example, the grant statement contains the permissions required by IBM MQ classes for JMS. To use these grant statements in your policy configuration file, you might need to modify the path names depending on where you have installed IBM MQ classes for JMS and where you store your applications.

The sample applications supplied with IBM MQ classes for JMS, and scripts to run them, do not enable the security manager.

## ClassName whitelisting in JMS ObjectMessage: **V 9.0.1**

In IBM MQ classes for JMS, support for whitelisting of classes in the implementation of the JMS ObjectMessage interface provides a potential mitigation against some of the security risks that potentially relate to the Java object serialization and deserialization mechanism.

The Java object serialization and deserialization mechanism has been identified as a potential security risk because deserialization instantiates arbitrary Java objects, where there is the potential for maliciously sent data to cause various problems. One notable application of serialization is in Java Message Service (JMS) ObjectMessages that use serialization to encapsulate and transfer arbitrary objects.

Serialization whitelisting is a potential mitigation against some of the risks that serialization poses. By explicitly specifying which classes can be encapsulated in, and extracted from, ObjectMessages, whitelisting provides some protection against some serialization risks.

### Whitelisting in IBM MQ classes for JMS

From IBM MQ Version 9.0.1 Continuous Delivery release, IBM MQ classes for JMS supports whitelisting of classes in the implementation of the JMS ObjectMessage interface.

**V 9.0.0.1** With APAR IT14385, and from IBM MQ Version 9.0.0, Fix Pack 1, IBM MQ classes for JMS supports whitelisting of classes in the implementation of the JMS ObjectMessage interface in the IBM MQ Version 9.0.0 Long Term Support release.

The whitelist defines which Java classes might be serialized with ObjectMessage.setObject() and deserialized with ObjectMessage.getObject().

Attempts to serialize or deserialize an instance of a class not included in the whitelist with ObjectMessage cause a javax.jms.MessageFormatException to be thrown, with a java.io.InvalidClassException as its cause.

### Producing the whitelist

**Important:** IBM MQ classes for JMS cannot be distributed with a whitelist. The choice of classes to be transferred by using ObjectMessages is an application design choice and IBM MQ cannot pre-empt that.

For this reason, the whitelisting mechanism allows for two modes of operation:

#### DISCOVERY

In this mode, the mechanism produces a listing of fully qualified class names, reporting all classes that have been observed to be serialized or deserialized in ObjectMessages.

#### ENFORCEMENT

In this mode, the mechanism enforces whitelisting, rejecting attempts to serialize or deserialize classes that are not in the whitelist.

If you want to use this mechanism, you must initially run in DISCOVERY mode to gather the list of currently serialized and deserialized classes, review the list and use it as a basis for your whitelist. It might even be appropriate to use the list unchanged, but the list must be reviewed first before you decide to do this.

### Controlling the whitelisting mechanism

Three system properties are available to control the whitelisting mechanism:

#### **com.ibm.mq.jms.whitelist**

This property can be specified in either of the following ways:

- The path name of the file that contains the whitelist, in file URI format (that is, starting with file:). In DISCOVERY mode, this file is written to by the whitelisting mechanism. The file must not exist. If the file does exist, the mechanism throws an exception rather than overwrite it. In ENFORCEMENT mode, this file is read by the whitelisting mechanism.
- A comma-separated list of fully qualified class names that constitute the whitelist.

If this property is unset, the whitelist mechanism is inactive.

If you are using a Java Security Manager, you must ensure that the IBM MQ classes for JMS JAR files have read and write access to this file.

#### **com.ibm.mq.jms.whitelist.discover**

- If this property is unset or set to false, the whitelist mechanism runs in ENFORCEMENT mode.
- If this property is set to true and the whitelist has been specified as a file URI, the whitelist mechanism runs in DISCOVERY mode.
- If this property is set to true and the whitelist has been specified as a list of class names, the whitelist mechanism throws a suitable exception.
- If this property is set to true and the whitelist has not been specified by using the com.ibm.mq.jms.whitelist property, the whitelist mechanism is inactive.
- If this property is set to true and the whitelist file already exists, the whitelist mechanism throws a java.io.InvalidClassException and entries are not added to the file.

#### **com.ibm.mq.jms.whitelist.mode**

This string property can be specified in any of three ways:

- If this property is set to SERIALIZE, then ENFORCEMENT mode performs whitelist validation only on the ObjectMessage.setObject() method.
- If this property is set to DESERIALIZE, then ENFORCEMENT mode performs whitelist validation only on the ObjectMessage.getObject() method.
- If this property is unset, or set to any other value, then ENFORCEMENT mode performs whitelist validation on both the ObjectMessage.getObject() and the ObjectMessage.setObject() methods.



#### **Format of the whitelist file**

These are the main features of the format of the whitelist file:

- The whitelist file is in default platform file encoding with platform-appropriate line-endings.

**Note:** If a whitelist file is being used, then that file is always written and read using the default file encoding for the JVM.

This is fine if the whitelist file is generated in any of the following ways:

-  Generated by a stand-alone application running on z/OS and used by other stand-alone applications that are also running on z/OS.
- Generated by an application running inside of WebSphere Application Server on any platform, and used by another instance of WebSphere Application Server.
-  Generated by a stand-alone application running on IBM MQ for Multiplatforms, and used by other stand-alone applications running on IBM MQ for Multiplatforms, or by applications running inside of WebSphere Application Server on any platform.

However, as WebSphere Application Server uses ASCII, and a standalone JVM uses EBCDIC, there will be file encoding issues if the whitelist file is generated in either of the following ways:

- Generated on z/OS, then used by standalone applications running on a platform other than z/OS or by WebSphere Application Server.



- Generated by either WebSphere Application Server or a standalone application running on a platform other than z/OS, then used by a stand-alone application on z/OS.
- Each non-empty line contains a fully qualified class name. Empty lines are ignored.
- Comments can be included - anything following a '#' character, to the end of the line, is ignored.
- There is a very basic wildcarding mechanism:
  - '\*' can be the **last** element of a class name.
  - '\*' matches a **single** element of a class name, that is, the class, but no part of the package.
 So `com.ibm.mq.*` would match `com.ibm.mq.MQMessage` but not `com.ibm.mq.jmqi.remote.api.RemoteFAP`.
- Wildcarding does not work for classes in the default package that is for classes without an explicit package name, so a class name of "\*" is rejected.
- Badly formatted whitelist files, for example, files that contain an entry such as `com.ibm.mq.*.Message`, where the wildcard is not the last element, result in a `java.lang.IllegalArgumentException` being thrown.
- An empty whitelist file has the effect of totally disabling the use of `ObjectMessage`.

### Format of the whitelist as a comma-separated list

The same wildcarding mechanism is available for a whitelist as a comma-separated list.

- The '\*' can be expanded by the operating system if specified on a command line or in a shell script or batch file, so it might need special handling.
- The '#' comment character is only applicable when a file is specified. If the whitelist is specified as a comma-separated list of class names, then assuming that the operating system or shell doesn't process it, as it is the default comment character in many UNIX or Linux shells, it is treated as a normal character.

### When does whitelisting happen?

Whitelisting is initiated when the application first runs an `ObjectMessage setMessage()` or `getMessage()` method.

The system properties are evaluated, the whitelist file is opened and in ENFORCEMENT mode, the list of whitelisted classes are loaded when the mechanism is initialized. At this point, an entry is written in to the IBM MQ JMS log file for the application.

When the mechanism is initialized, its parameters might not be changed. As the time of initialization is not easily predicted as it depends on application behavior. The system property settings and the whitelist file contents should therefore be regarded as fixed from the time that the application is started. Do not change the properties or the contents of the whitelist file while the application is running, as the results are not guaranteed.

### Points to consider

The best approach to mitigating the risks intrinsic to the Java serialization mechanism would be to explore alternative approaches to data transfer such as using JSON instead of `ObjectMessage`. Using Advanced Message Security (AMS) mechanisms can add further security by ensuring that messages come from trusted sources.

If you use the Java Security Manager mechanism with your application, you must grant the following permissions:

- `FilePermission` on any whitelist file that you use, with read permission for ENFORCEMENT mode, write permission for DISCOVER mode.
- `PropertyPermission` (read) on the `com.ibm.mq.jms.whitelist`, `com.ibm.mq.jms.whitelist.discover`, and `com.ibm.mq.jms.whitelist.mode` properties.

### Related concepts:

“Running IBM MQ classes for JMS applications under the Java Security Manager” on page 845  
IBM MQ classes for JMS can run with the Java security manager enabled. To run applications successfully with the Java Security Manager enabled, you must configure your Java virtual machine (JVM) with a suitable policy configuration file.

### Post installation setup for IBM MQ classes for JMS applications:

This topic tells you what authorities IBM MQ classes for JMS applications need in order to access the resources of a queue manager. It also introduces connection modes and describes how to configure a queue manager so that applications can connect in client mode.

**Remember to check the IBM MQ readme file. It might contain information that supersedes the information in this topic.**

*Objects used by JMS that require authorization for non-privileged users:*

Non-privileged users need authorization granted to access the queues used by JMS. Every JMS application needs authorization to the queue manager with which it works.

For details about access control in IBM MQ, see [Setting up security](#).

IBM MQ classes for JMS applications need connect and inq authority to the queue manager. You can set appropriate authorizations using the **setmqaut** control command, for example:

```
setmqaut -m QM1 -t qmgr -g jmsappsgroup +connect +inq
```

For the point-to-point domain, the following authorities are required:

- Queues that are used by MessageProducer objects need put authority.
- Queues that are used by MessageConsumer and QueueBrowser objects need get, inq, and browse authorities.
- The QueueSession.createTemporaryQueue() method needs access to the model queue specified by the TEMPMODEL property of the QueueConnectionFactory object. By default this model queue is SYSTEM.TEMP.MODEL.QUEUE.

If any of these queues are alias queues, their target queues require inquire authority. If the target queue is a cluster queue it also requires browse authority.

For the publish/subscribe domain, the following queues are used if the IBM MQ classes for JMS are connecting to an IBM MQ queue manager in IBM MQ messaging provider migration mode:

- SYSTEM.JMS.ADMIN.QUEUE
- SYSTEM.JMS.REPORT.QUEUE
- SYSTEM.JMS.MODEL.QUEUE
- SYSTEM.JMS.PS.STATUS.QUEUE
- SYSTEM.JMS.ND.SUBSCRIBER.QUEUE
- SYSTEM.JMS.D.SUBSCRIBER.QUEUE
- SYSTEM.JMS.ND.CC.SUBSCRIBER.QUEUE
- SYSTEM.JMS.D.CC.SUBSCRIBER.QUEUE
- SYSTEM.BROKER.CONTROL.QUEUE

For further information on IBM MQ messaging provider migration mode, see [Configuring the JMS PROVIDERVERSION](#) property

Additionally, if the IBM MQ classes for JMS are connecting to a queue manager in this mode, any application that publishes messages needs access to the stream queue specified by the `TopicConnectionFactory` or topic object. By default, this queue is `SYSTEM.BROKER.DEFAULT.STREAM`.

If you use `ConnectionConsumer`, IBM MQ Resource Adapter, or the WebSphere Application Server IBM MQ messaging provider, additional authorization might be needed.

Queues to be read by the `ConnectionConsumer` must have `get`, `inq`, and `browse` authorities. The system dead-letter queue, and any backout-requeue queue or report queue used by the `ConnectionConsumer` must have `put` and `passall` authorities.

When an application uses IBM MQ messaging provider normal mode to perform publish/subscribe messaging, the application makes use of the integrated publish/subscribe functionality provided by the queue manager. See *Publish/subscribe security* for information on securing the topics and queues that are used.

*Connection modes for IBM MQ classes for JMS:*

An IBM MQ classes for JMS application can connect to a queue manager in either client or bindings mode. In client mode, IBM MQ classes for JMS connects to the queue manager over TCP/IP. In bindings mode, IBM MQ classes for JMS connects directly to the queue manager using the Java Native Interface (JNI).

An application running in WebSphere Application Server on z/OS can connect to a queue manager in either bindings or client mode, but an application running in any other environment on z/OS can connect to a queue manager only in bindings mode. An application running on any other platform can connect to a queue manager in either bindings or client mode.

You can use the current or any earlier supported version of IBM MQ classes for JMS with a current queue manager, and you can use a current or earlier supported version of queue manager with the current version of IBM MQ classes for JMS. If you mix different versions, function is limited to the level of the earlier version.

The following sections describe each of the connection modes in more detail.

### **Client mode**

To connect to a queue manager in client mode, an IBM MQ classes for JMS application can run on the same system on which the queue manager is running, or on a different system. In each case, IBM MQ classes for JMS connects to the queue manager over TCP/IP.

### **Bindings mode**

To connect to a queue manager in bindings mode, an IBM MQ classes for JMS application must run on the same system on which the queue manager is running.

The IBM MQ classes for JMS connects directly to the queue manager using the Java Native Interface (JNI). To use the bindings transport, the IBM MQ classes for JMS must be run in an environment that has access to the IBM MQ Java Native Interface libraries; see “Configuring the Java Native Interface (JNI) libraries” on page 835 for further information.

The IBM MQ classes for JMS support the following values for *ConnectOption*:

- `MQCNO_FASTPATH_BINDING`
- `MQCNO_STANDARD_BINDING`
- `MQCNO_SHARED_BINDING`

- MQCNO\_ISOLATED\_BINDING
- MQCNO\_SERIALIZE\_CONN\_TAG\_QSG
- MQCNO\_RESTRICT\_CONN\_TAG\_QSG
- MQCNO\_SERIALIZE\_CONN\_TAG\_Q\_MGR
- MQCNO\_RESTRICT\_CONN\_TAG\_Q\_MGR

To change the connection options used by the IBM MQ classes for JMS, modify the Connection Factory property CONNOPT.

For further information on connection options, see “Connecting to a queue manager using the MQCONN call” on page 94

To use the bindings transport, the Java Runtime Environment being used must support the Coded Character Set Identifier (CCSID) of the queue manager that the IBM MQ classes for JMS are connecting to.

Details on how to determine what CCSIDs are supported by a Java Runtime Environment can be found in IBM MQ FDC with Probe ID 21 generated when using the IBM MQ V7 classes for Java or IBM MQ V7 classes for JMS.

*Configuring your queue manager so that IBM MQ classes for JMS applications can connect in client mode:*

To configure your queue manager so that IBM MQ classes for JMS applications can connect in client mode, you must create a server connection channel definition and start a listener.

### Creating a server connection channel definition

On all platforms, you can use the MQSC command DEFINE CHANNEL to create a server connection channel definition. See the following example:

```
DEFINE CHANNEL(JAVA.CHANNEL) CHLTYPE(SVRCONN) TRPTYPE(TCP)
```

**IBM i** On IBM i, you can use the CL command CRTMQMCHL instead, as in the following example:

```
CRTMQMCHL CHLNAME(JAVA.CHANNEL) CHLTYPE(*SVRCN)
TRPTYPE(*TCP)
MQMNAME(QMGRNAME)
```

In this command, *QMGRNAME* is the name of your queue manager.

**Linux** **Windows** On Linux and Windows, you can also create a server connection channel definition using IBM MQ Explorer.

**z/OS** On z/OS you can use the operations and control panels to create a server connection channel definition .

The name of the channel (JAVA.CHANNEL in the previous examples) must be the same as the channel name specified by the CHANNEL property of the connection factory that your application uses to connect to the queue manager. The default value of the CHANNEL property is SYSTEM.DEF.SVRCONN.

### Starting a listener

You must start a listener for your queue manager if one is not already started.

**Multi**

On Multiplatforms, you can use the MQSC command `START LISTENER` to start a listener after first creating a listener object by using the MQSC command `DEFINE LISTENER`, as shown in the following example:

```
DEFINE LISTENER(LISTENER.TCP) TRPTYPE(TCP) PORT(1414)
START LISTENER(LISTENER.TCP)
```

**z/OS**

On z/OS, you use only the `START LISTENER` command, as in the following example, but note that the channel initiator address space must be started before you can start a listener:

```
START LISTENER TRPTYPE(TCP) PORT(1414)
```

**IBM i**

On IBM i, you can also use the CL command `STRMQMLSR` to start a listener, as in the following example:

```
STRMQMLSR PORT(1414) QMGRNAME(QMGRNAME)
```

In this command, *QMGRNAME* is the name of your queue manager.

**ULW**

On UNIX, Linux, and Windows, you can also use the control command `runmq1sr` to start a listener, as in the following example:

```
runmq1sr -t tcp -p 1414 -m QMgrName
```

In this command, *QMgrName* is the name of your queue manager.

**Linux****Windows**

On Linux and Windows, you can also start a listener using IBM MQ Explorer.

**z/OS**

On z/OS, you can also use the operations and control panels to start a listener.

The number of the port on which the listener is listening must be the same as the port number specified by the `PORT` property of the connection factory that your application uses to connect to the queue manager. The default value of the `PORT` property is 1414.

### The point-to-point IVT for IBM MQ classes for JMS:

A point-to-point installation verification test (IVT) program is supplied with IBM MQ classes for JMS. The program connects to a queue manager in either bindings or client mode, sends a message to the queue called `SYSTEM.DEFAULT.LOCAL.QUEUE`, and then receives the message from the queue. The program can create and configure all the objects that it requires dynamically at run time, or it can use JNDI to retrieve administered objects from a directory service.

Run the installation verification test without using JNDI first because the test is self contained and does not require the use of a directory service. For a description of administered objects, see *Configuring JMS objects using the administration tool*.

### The point-to-point installation verification test without using JNDI

In this test, the IVT program creates and configures all the objects that it requires dynamically at run time and does not use JNDI.

A script is provided to run the IVT program. The script is called `IVTRun` on UNIX and Linux systems and `IVTRun.bat` on Windows, and is in the `bin` subdirectory of the IBM MQ classes for JMS installation directory.

To run the test in bindings mode, enter the following command:

```
IVTRun -nojndi [-m qmgr] [-v providerVersion] [-t]
```

To run the test in client mode, first set up the queue manager as described in “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470 noting that the channel to be used defaults to SYSTEM.DEF.SVRCONN and the queue to be used is SYSTEM.DEFAULT.LOCAL.QUEUE, then enter the following command:

```
IVTRun -nojndi -client -m qmgr -host hostname [-port port] [-channel channel]
[-v providerVersion] [-ccsid ccsid] [-t]
```

No equivalent script is provided on z/OS systems, but you can run the IVT in bindings mode by invoking the Java class directly, using the following command:

```
java com.ibm.mq.jms.MQJMSIVT -nojndi [-m qmgr] [-v providerVersion] [-t]
```

The classpath must contain com.ibm.mqjms.jar.

The parameters on the commands have the following meanings:

**-m *qmgr***

The name of the queue manager to which the IVT program connects. If you run the test in bindings mode and omit this parameter, the IVT program connects to the default queue manager.

**-host *hostname***

The host name or IP address of the system on which the queue manager is running.

**-port *port***

The number of the port on which the listener of the queue manager is listening. The default value is 1414.

**-channel *channel***

The name of the MQI channel that the IVT program uses to connect to the queue manager. The default value is SYSTEM.DEF.SVRCONN.

**-v *providerVersion***

The release level of the queue manager to which the IVT program expects to connect.

This parameter is used to set the PROVIDERVERSION property of an MQQueueConnectionFactory object and has the same valid values as those of the PROVIDERVERSION property. For more information about this parameter therefore, including its valid values, see JMS **PROVIDERVERSION** property and the description of the PROVIDERVERSION property in Properties of IBM MQ classes for JMS objects.

The default value is unspecified.

**-ccsid *ccsid***

The identifier (CCSID) of the coded character set, or code page, to be used by the connection. The default value is 819.

**-t** Tracing is enabled. By default, tracing is disabled.

A successful test produces output similar to the following sample output:

```
5724-H72, 5655-R36, 5724-L26, 5655-L82 (c) Copyright IBM Corp. 2008. All Rights Reserved.
WebSphere MQ classes for Java(tm) Message Service 7.0
Installation Verification Test
```

```
Creating a QueueConnectionFactory
Creating a Connection
Creating a Session
Creating a Queue
Creating a QueueSender
Creating a QueueReceiver
Creating a TextMessage
```

Sending the message to SYSTEM.DEFAULT.LOCAL.QUEUE  
Reading the message back again

```
Got message
JMSMessage class: jms_text
JMSType: null
JMSDeliveryMode: 2
JMSExpiration: 0
JMSPriority: 4
JMSMessageID: ID:414d5120514d5f6d6277202020202001edb14620005e03
JMSTimestamp: 1187170264000
JMSCorrelationID: null
JMSDestination: queue:///SYSTEM.DEFAULT.LOCAL.QUEUE
JMSReplyTo: null
JMSRedelivered: false
JMSXUserID: mwhite
JMS_IBM_Encoding: 273
JMS_IBM_PutApplType: 28
JMSXAppID: IBM MQ Client for Java
JMSXDeliveryCount: 1
JMS_IBM_PutDate: 20070815
JMS_IBM_PutTime: 09310400
JMS_IBM_Format: MQSTR
JMS_IBM_MsgType: 8
A simple text message from the MQJMSIVT
Reply string equals original string
Closing QueueReceiver
Closing QueueSender
Closing Session
Closing Connection
IVT completed OK
IVT finished
```

### The point-to-point installation verification test using JNDI

In this test, the IVT program uses JNDI to retrieve administered objects from a directory service.

Before you can run the test, you must configure a directory service that is based on a Lightweight Directory Access Protocol (LDAP) server or the local file system. You must also configure the IBM MQ JMS administration tool so that it can use the directory service to store administered objects. For more information about these prerequisites, see “Prerequisites for IBM MQ classes for JMS” on page 827. For information about how to configure the IBM MQ JMS administration tool, see Configuring the JMS administration tool.

The IVT program must be able to use JNDI to retrieve an MQQueueConnectionFactory object and an MQQueue object from the directory service. A script is provided to create these administered objects for you. The script is called IVTSetup on UNIX and Linux systems and IVTSetup.bat on Windows, and is in the bin subdirectory of the IBM MQ classes for JMS installation directory. To run the script, enter the following command:

```
IVTSetup
```

The script invokes the IBM MQ JMS administration tool to create the administered objects.

The MQQueueConnectionFactory object is bound with the name ivtQCF and is created with the default values for all its properties, which means that the IVT program runs in bindings mode and connects to the default queue manager. If you want the IVT program to run in client mode, or connect to a queue manager other than the default queue manager, you must use the IBM MQ JMS administration tool or IBM MQ Explorer to change the appropriate properties of the MQQueueConnectionFactory object. For

information about how to use the IBM MQ Explorer JMS administration tool, see Configuring JMS objects using the administration tool. For information about how to use IBM MQ Explorer, see Introduction to IBM MQ Explorer or the help provided with IBM MQ Explorer.

The MQQueue object is bound with the name `ivtQ` and is created with the default values for all its properties, except for the `QUEUE` property, which has the value `SYSTEM.DEFAULT.LOCAL.QUEUE`.

When you have created the administered objects, you can run the IVT program. To run the test using JNDI, enter the following command:

```
IVTRun -url "providerURL" [-icf initCtxFact] [-t]
```

The parameters on the command have the following meanings:

**-url "*providerURL*"**

The uniform resource locator (URL) of the directory service. The URL can have one of the following formats:

- `ldap://hostname/contextName` , for a directory service based on an LDAP server
- `file:/directoryPath` , for a directory service based on the local file system

Note that you must enclose the URL in quotation marks (").

**-icf *initCtxFact***

The class name of the initial context factory, which must be one of the following values:

- `com.sun.jndi.ldap.LdapCtxFactory`, for a directory service based on an LDAP server. This is the default value.
- `com.sun.jndi.fscontext.RefFSContextFactory`, for a directory service based on the local file system.

**-t** Tracing is enabled. By default, tracing is disabled.

A successful test produces output similar to that for a successful test without using JNDI. The main difference is that the output indicates that the test is using JNDI to retrieve an `MQQueueConnectionFactory` object and an `MQQueue` object.

Although not strictly necessary, it is good practice to tidy up after the test by deleting the administered objects created by the `IVTSetup` script. A script is provided for this purpose. The script is called `IVTTidy` on UNIX and Linux systems and `IVTTidy.bat` on Windows, and is in the `bin` subdirectory of the IBM MQ classes for JMS installation directory.

### **Problem determination for the point-to-point installation verification test**

The installation verification test might fail for the following reasons:

- If the IVT program writes a message indicating that it cannot find a class, check that your class path is set correctly, as described in "Setting environment variables" on page 833.
- The test might fail with the following message:

```
Failed to connect to queue manager ' qmgr ' with connection mode ' connMode '
and host name ' hostname '
```

and an associated reason code of 2059. The variables in the message have the following meanings:

***qmgr*** The name of the queue manager to which the IVT program is trying to connect. This message insert is blank if the IVT program is trying to connect to the default queue manager in bindings mode.

***connMode***

The connection mode, which is either Bindings or Client.



### *hostname*

The host name or IP address of the system on which the queue manager is running.

This message means that the queue manager to which the IVT program is trying to connect is not available. Check that the queue manager is running and, if the IVT program is trying to connect to the default queue manager, make sure that the queue manager is defined as the default queue manager for your system.

- The test might fail with the following message:

```
Failed to open MQ queue 'SYSTEM.DEFAULT.LOCAL.QUEUE'
```

This message means that the queue SYSTEM.DEFAULT.LOCAL.QUEUE does not exist on the queue manager to which the IVT program is connected. Alternatively, if the queue does exist, the IVT program cannot open the queue because it is not enabled for putting and getting messages. Check that the queue exists and that it is enabled for putting and getting messages.

- The test might fail with the following message:

```
Unable to bind to object
```

This message means that there is a connection to the LDAP server, but that the LDAP server is not correctly configured. Either the LDAP server is not configured for storing Java objects, or the permissions on the objects or the suffix are not correct. For more help in this situation, see the documentation for your LDAP server.

- The test might fail with the following message:

```
The security authentication was not valid that was supplied for
QueueManager ' qmgr ' with connection mode 'Client' and host name ' hostname '
```

This message means that the queue manager is not correctly set up to accept a client connection from your system. See “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470 for details.

### **The publish/subscribe IVT for IBM MQ classes for JMS:**

A publish/subscribe installation verification test (IVT) program is supplied with IBM MQ classes for JMS. The program connects to a queue manager in either bindings or client mode, subscribes to a topic, publishes a message on the topic, and then receives the message that it has just published. The program can create and configure all the objects that it requires dynamically at run time, or it can use JNDI to retrieve administered objects from a directory service.

Run the installation verification test without using JNDI first because the test is self contained and does not require the use of a directory service. For a description of administered objects, see Configuring JMS objects using the administration tool.

### **The publish/subscribe installation verification test without using JNDI**

In this test, the IVT program creates and configures all the objects that it requires dynamically at run time and does not use JNDI.

A script is provided to run the IVT program. The script is called PSIVTRun on UNIX and Linux systems and PSIVTRun.bat on Windows, and is in the bin subdirectory of the IBM MQ classes for JMS installation directory.

To run the test in bindings mode, enter the following command:

```
PSIVTRun -nojndi [-m qmgr] [-bqm brokerQmgr] [-v providerVersion] [-t]
```

To run the test in client mode, first set up the queue manager as described in “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470 noting that the channel to be used defaults to SYSTEM.DEF.SVRCONN, then enter the following command:

```
PSIVTRun -nojndi -client -m qmgr -host hostname [-port port] [-channel channel]
[-bqm brokerQmgr] [-v providerVersion] [-ccsid ccsid] [-t]
```

The parameters on the commands have the following meanings:

**-m *qmgr***

The name of the queue manager to which the IVT program connects. If you run the test in bindings mode and omit this parameter, the IVT program connects to the default queue manager.

**-host *hostname***

The host name or IP address of the system on which the queue manager is running.

**-port *port***

The number of the port on which the listener of the queue manager is listening. The default value is 1414.

**-channel *channel***

The name of the MQI channel that the IVT program uses to connect to the queue manager. The default value is SYSTEM.DEF.SVRCONN.

**-bqm *brokerQmgr***

The name of the queue manager on which the broker is running. The default value is the name of the queue manager to which the IVT program connects.

This parameter is not relevant for queue manager version number *v* of 7, or greater.

**-v *providerVersion***

The release level of the queue manager to which the IVT program expects to connect.

This parameter is used to set the PROVIDERVERSION property of an MQTopicConnectionFactory object and has the same valid values as those of the PROVIDERVERSION property. For more information about this parameter therefore, including its valid values, see the description of the PROVIDERVERSION property in Properties of IBM MQ classes for JMS objects.

The default value is unspecified.

**-ccsid *ccsid***

The identifier (CCSID) of the coded character set, or code page, to be used by the connection. The default value is 819.

**-t** Tracing is enabled. By default, tracing is disabled.

A successful test produces output similar to the following sample output:

```
5724-H72, 5655-R36, 5724-L26, 5655-L82 (c) Copyright IBM Corp. 2008. All Rights Reserved.
IBM MQ classes for Java(tm) Message Service 7.0
Publish/Subscribe Installation Verification Test
```

```
Creating a TopicConnectionFactory
Creating a Connection
Creating a Session
Creating a Topic
Creating a TopicPublisher
Creating a TopicSubscriber
Creating a TextMessage
Adding text
Publishing the message to topic://MQJMS/PSIVT/Information
Waiting for a message to arrive [5 secs max]...
```

```
Got message:
JMSMessage class: jms_text
JMSType: null
JMSDeliveryMode: 2
JMSExpiration: 0
JMSPriority: 4
JMSMessageID: ID:414d5120514d5f6d6277202020202001edb14620006706
JMSTimestamp: 1187182520203
JMSCorrelationID: ID:414d5120514d5f6d6277202020202001edb14620006704
JMSDestination: topic://MQJMS/PSIVT/Information
JMSReplyTo: null
JMSRedelivered: false
JMSXUserID: mwhite
JMS_IBM_Encoding: 273
JMS_IBM_PutAppType: 26
JMSXAppID: QM_mbw
JMSXDeliveryCount: 1
JMS_IBM_PutDate: 20070815
JMS_IBM_ConnectionID: 414D5143514D5F6D6277202020202001EDB14620006601
JMS_IBM_PutTime: 12552020
JMS_IBM_Format: MQSTR
JMS_IBM_MsgType: 8
A simple text message from the MQJMSPSIVT program
Reply string equals original string
Closing TopicSubscriber
Closing TopicPublisher
Closing Session
Closing Connection
PSIVT finished
```

### **The publish/subscribe installation verification test using JNDI**

In this test, the IVT program uses JNDI to retrieve administered objects from a directory service.

Before you can run the test, you must configure a directory service that is based on a Lightweight Directory Access Protocol (LDAP) server or the local file system. You must also configure the IBM MQ JMS administration tool so that it can use the directory service to store administered objects. For more information about these prerequisites, see “Prerequisites for IBM MQ classes for JMS” on page 827. For information about how to configure the IBM MQ JMS administration tool, see Configuring the JMS administration tool.

The IVT program must be able to use JNDI to retrieve an MQTopicConnectionFactory object and an MQTopic object from the directory service. A script is provided to create these administered objects for you. The script is called IVTSetup on UNIX and Linux systems and IVTSetup.bat on Windows, and is in the bin subdirectory of the IBM MQ classes for JMS installation directory. To run the script, enter the following command:

```
IVTSetup
```

The script invokes the IBM MQ JMS administration tool to create the administered objects.

The MQTopicConnectionFactory object is bound with the name ivtTCF and is created with the default values for all its properties, which means that the IVT program runs in bindings mode, connects to the default queue manager, and uses the embedded publish/subscribe function. If you want the IVT program to run in client mode, connect to a queue manager other than the default queue manager, or use IBM Integration Bus instead of the embedded publish/subscribe function, you must use the IBM MQ JMS administration tool or IBM MQ Explorer to change the appropriate properties of the MQTopicConnectionFactory object. For information about how to use the IBM MQ JMS administration

tool, see Configuring JMS objects using the administration tool. For information about how to use IBM MQ Explorer, see the help provided with IBM MQ Explorer.

The MQTopic object is bound with the name ivtT and is created with the default values for all its properties, except for the TOPIC property, which has the value MQJMS/PSIVT/Information.

When you have created the administered objects, you can run the IVT program. To run the test using JNDI, enter the following command:

```
PSIVTRun -url "providerURL" [-icf initCtxFact] [-t]
```

The parameters on the command have the following meanings:

**-url "providerURL"**

The uniform resource locator (URL) of the directory service. The URL can have one of the following formats:

- `ldap://hostname/contextName` , for a directory service based on an LDAP server
- `file:/directoryPath` , for a directory service based on the local file system

Note that you must enclose the URL in quotation marks (").

**-icf initCtxFact**

The class name of the initial context factory, which must be one of the following values:

- `com.sun.jndi.ldap.LdapCtxFactory`, for a directory service based on an LDAP server. This is the default value.
- `com.sun.jndi.fscontext.RefFSContextFactory`, for a directory service based on the local file system.

**-t** Tracing is enabled. By default, tracing is disabled.

A successful test produces output similar to that for a successful test without using JNDI. The main difference is that the output indicates that the test is using JNDI to retrieve an MQTopicConnectionFactory object and an MQTopic object.

Although not strictly necessary, it is good practice to tidy up after the test by deleting the administered objects created by the IVTSetup script. A script is provided for this purpose. The script is called IVTTidy on UNIX and Linux systems and IVTTidy.bat on Windows, and is in the bin subdirectory of the IBM MQ classes for JMS installation directory.

### Problem determination for the publish/subscribe installation verification test

The installation verification test might fail for the following reasons:

- If the IVT program writes a message indicating that it cannot find a class, check that your class path is set correctly, as described in "Setting environment variables" on page 833.
- The test might fail with the following message:

```
Failed to connect to queue manager ' qmgr ' with
connection mode ' connMode ' and host name ' hostname '
```

and an associated reason code of 2059. The variables in the message have the following meanings:

**qmgr** The name of the queue manager to which the IVT program is trying to connect. This message insert is blank if the IVT program is trying to connect to the default queue manager in bindings mode.

**connMode**

The connection mode, which is either Bindings or Client.

**hostname**

The host name or IP address of the system on which the queue manager is running.

This message means that the queue manager to which the IVT program is trying to connect is not available. Check that the queue manager is running and, if the IVT program is trying to connect to the default queue manager, make sure that the queue manager is defined as the default queue manager for your system.

- The test might fail with the following message:  
Unable to bind to object

This message means that there is a connection to the LDAP server, but that the LDAP server is not correctly configured. Either the LDAP server is not configured for storing Java objects, or the permissions on the objects or the suffix are not correct. For more help in this situation, see the documentation for your LDAP server.

- Th test might fail with the following message:  
The security authentication was not valid that was supplied for  
QueueManager ' qmgr ' with connection mode 'Client' and host name ' hostname '

This message means that the queue manager is not correct set up to accept a client connect from your system. See “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470 for details.

### Scripts provided with IBM MQ classes for JMS:

A number of scripts are provided to assist with common tasks that need to be performed when using IBM MQ classes for JMS.

Table 112 lists all the scripts and their uses. The scripts are in the bin subdirectory of the IBM MQ classes for JMS installation directory.

Table 112. Scripts provided with IBM MQ classes for JMS

Utility	Use
Cleanup <sup>1</sup>	This script is maintained for compatibility with previous releases but performs no function. Manual cleanup of subscription information is no longer necessary
DefaultConfiguration	Runs the default configuration application on platforms other than Windows.
formatLog <sup>1</sup>	This script is maintained for compatibility with previous releases but performs no function. Log output is now produced in readable text.
IVTRun <sup>1</sup> IVTSetup <sup>1</sup> IVTTidy <sup>1</sup>	Used in the point-to-point installation verification test, as described in “The point-to-point IVT for IBM MQ classes for JMS” on page 853.
JMSAdmin <sup>1</sup>	Runs the IBM MQ JMS administration tool, as described in Starting the administration tool.
JMSAdmin.config	The configuration file for the IBM MQ JMS administration tool, as described in Configuring the JMS administration tool.
PSIVTRun <sup>1</sup>	Runs the publish/subscribe installation verification test program, as described in “The publish/subscribe IVT for IBM MQ classes for JMS” on page 857.
PSReportDump.class	This class is maintained for compatibility with previous releases, but performs no function.
setjmsenv	Sets the environment variables for running an IBM MQ classes for JMS application in a 32-bit Java virtual machine (JVM) on UNIX and Linux systems, as described in “Setting environment variables” on page 833.

Table 112. Scripts provided with IBM MQ classes for JMS (continued)

Utility	Use
setjmsenv64	Sets the environment variables for running an IBM MQ classes for JMS application in a 64-bit JVM on UNIX and Linux systems, as described in "Setting environment variables" on page 833.

**Note:**

1. On Windows, the file name has the extension .bat .

**Support for OSGi:**

OSGi provides a framework that supports the deployment of applications as bundles. Nine OSGi bundles are supplied as part of the IBM MQ classes for JMS.

OSGi provides a general purpose, secure, and managed Java framework, which supports the deployment of applications that come in the form of bundles. OSGi-compliant devices can download and install bundles, and remove them when they are no longer required. The framework manages the installation and update of bundles in a dynamic and scalable fashion.

The IBM MQ classes for JMS includes the following OSGi bundles.

**com.ibm.msg.client.osgi.jmsversion\_number.jar**

The common layer of code in the IBM MQ classes for JMS. For information about the layered architecture of IBM MQ classes for JMS, see Layered architecture.

**com.ibm.msg.client.osgi.jms.prereq\_version\_number.jar**

The prerequisite Java archive (JAR) files for the common layer.

**com.ibm.msg.client.osgi.commonservices.j2se\_version\_number.jar**

Common services for Java Platform, Standard Edition (Java SE) applications.

**com.ibm.msg.client.osgi.nls\_version\_number.jar**

Messages for the common layer.

**com.ibm.msg.client.osgi.wmq\_version\_number.jar**

The IBM MQ messaging provider in IBM MQ classes for JMS. For information about the layered architecture of IBM MQ classes for JMS, see Layered architecture.

**com.ibm.msg.client.osgi.wmq.prereq\_version\_number.jar**

The prerequisite JAR files for the IBM MQ messaging provider.

**com.ibm.msg.client.osgi.wmq.nls\_version\_number.jar**

Messages for the IBM MQ messaging provider.

**com.ibm.mq.osgi.allclient\_version\_number.jar**

This JAR file allows applications to use both the IBM MQ classes for JMS and the IBM MQ classes for Java, and also includes the code to handle PCF messages.

**com.ibm.mq.osgi.allclientprereqs\_version\_number.jar**

This JAR file provides the prerequisites for com.ibm.mq.osgi.allclient\_version\_number.jar where *version\_number* is the version number of IBM MQ that is installed.

The bundles are installed into the java/lib/OSGi subdirectory of your IBM MQ installation, or the java\lib\OSGi folder on Windows.

From IBM MQ Version 8.0, use the bundles com.ibm.mq.osgi.allclient\_8.0.0.0.jar, and com.ibm.mq.osgi.allclientprereqs\_8.0.0.0.jar for any new applications. Using these bundles removes the restriction of not being able to run both IBM MQ classes for JMS and the IBM MQ classes for Java

within the same OSGi framework, all other restrictions still apply however. For versions of the product before IBM MQ Version 8.0, this restriction of using either IBM MQ classes for JMS or IBM MQ classes for Java applies.

The bundle `com.ibm.mq.osgi.javaversion_number.jar`, which is also installed into the `java/lib/OSGi` subdirectory of your IBM MQ installation, or the `java\lib\OSGi` folder on Windows, is part of the IBM MQ classes for Java. This bundle must not be loaded into an OSGi runtime environment that has the IBM MQ classes for JMS loaded.

The OSGi bundles for the IBM MQ classes for JMS have been written to the OSGi Release 4 specification. They do not work in an OSGi Release 3 environment.

You must set your system path or library path correctly so that the OSGi runtime environment can find any required DLL files or shared libraries.

If you use the OSGi bundles for the IBM MQ classes for JMS, temporary topics do not work. In addition, channel exit classes written in Java are not supported because of an inherent problem in loading classes in a multiple class loader environment such as OSGi. A user bundle can be aware of the IBM MQ classes for JMS bundles, but the IBM MQ classes for JMS bundles are not aware of any user bundle. As a result, the class loader used in an IBM MQ classes for JMS bundle cannot load a channel exit class that is in a user bundle.

For more information about OSGi, see the OSGi Alliance website.

### **Solving problems with IBM MQ classes for JMS:**

You can investigate problems by running the installation verification programs, and by using the trace and log facilities.

If a program does not complete successfully, run one of the installation verification programs, as described in “The point-to-point IVT for IBM MQ classes for JMS” on page 853 and “The publish/subscribe IVT for IBM MQ classes for JMS” on page 857, and follow the advice given in the diagnostic messages.

#### **Related information:**

Tracing IBM MQ classes for JMS applications

Tracing using `MQJMS_TRACE_LEVEL`

*Logging and IBM MQ classes for JMS:*

By default, log output is sent to the `mjqms.log` file. You can redirect it to a specific file or directory.

The IBM MQ classes for JMS log facility is provided to report serious problems, particularly problems that might indicate configuration errors rather than programming errors. By default, log output is sent to the `mjqms.log` file in the JVM working directory.

You can redirect log output to another file by setting the property `com.ibm.msg.client.commonservices.log.outputName`. The value for this property can be:

- A single path name.
- A comma-separated list of path names (all data is logged to all files).

Each path name can be:

- Absolute or relative.
- `stderr` or `System.err` to represent the standard error stream.
- `stdout` or `System.out` to represent the standard output stream.

If the value of the property identifies a directory, log output is written to `mqjms.log` in that directory. If the value of the property identifies a specific file, log output is written to that file.

You can set this property in the IBM MQ classes for JMS configuration file or as a system property on the **java** command. In the following example, the property is set as a system property and identifies a specific file:

```
java -Djava.library.path= library_path
-Dcom.ibm.msg.client.commonservices.log.outputName=/mydir/mylog.txt
MyAppClass
```

In the command, *library\_path* is the path to the directory containing the IBM MQ classes for JMS libraries (see “Configuring the Java Native Interface (JNI) libraries” on page 835 ).

You can disable log output by setting the property `com.ibm.msg.client.commonservices.log.status` to OFF. The default value of this property is ON.

The values `System.err` and `System.out` can be set to send log output to the `System.err` and `System.out` streams.

## Obtaining the IBM MQ classes for JMS separately

The IBM MQ classes for JMS are available within a self-extracting JAR file that you can download from Fix Central if you want to obtain just the IBM MQ classes for JMS JAR files, for deployment into a software management tool, or to use with stand-alone client applications.

### Before you begin

Before starting this task, make sure that you have a Java runtime environment (JRE) installed on your machine and that the JRE has been added to the system path.

The Java installer that is used in this installation process does not require running as root or any specific user. The only requirement is that the user it is run as has access write to the directory that you want the files to go in.

### About this task

Before IBM MQ Version 8.0, the IBM WebSphere MQ classes for Java or IBM WebSphere MQ classes for JMS are not available as a separate download. For IBM WebSphere MQ Version 7.5 or earlier, if you are developing and running Java language applications that use either the IBM WebSphere MQ classes for Java or IBM WebSphere MQ classes for JMS, you need to install them either by performing a full server installation or by installing one of the client SupportPacs onto the system where the application is being developed and the system where the application will run. This installation installs many more files than the IBM WebSphere MQ classes for Java and IBM WebSphere MQ classes for JMS files.

However, from IBM MQ Version 8.0, the following files are available within a self-extracting JAR file, which minimizes the size of the download and installation, and the time that is required to perform the installation:

- The IBM MQ classes for JMS
- The IBM MQ classes for Java
- The IBM MQ resource adapter
- The IBM MQ OSGi bundles

When you run the executable JAR file, it displays the IBM MQ license agreement, which must be accepted. It asks for a directory in which to install the IBM MQ classes for Java, IBM MQ classes for JMS, the resource adapter, and OSGi bundles. If the selected installation directory does not exist, it is created and the program files are installed. However, if the directory exists, an error is reported and no files are installed.



## Procedure

1. Download the IBM MQ Java JAR file from Fix Central. To locate the latest version that is available for download, enter the phrase "Java" in the **Text Search** box. The name of the file to be downloaded is in the format *V.R.M.F*-WS-MQ-Install-Java-All.jar where *V.R.M.F* is the product version number, for example 9.0.0.0.

If you cannot find the file, make sure that **Product Selected** is WebSphere MQ and **Version** is 9.0.

2. Start the installation from the directory to which you downloaded the file. To start the installation, enter a command in the following format:

```
java -jar V.R.M.F-WS-MQ-Install-Java-All.jar
```

where *V.R.M.F* is the product version number, for example 9.0.0.0, and *V.R.M.F*-WS-MQ-Install-Java-All.jar is the name of the file that was downloaded from Fix Central. For example, to install the IBM MQ classes for JMS for the Version 9.0.0.0 release, you would use the following command:

```
java -jar 9.0.0.0-WS-MQ-Install-Java-All.jar
```

**Note:** To carry out this installation, you must have a JRE installed on your machine and added to the system path.

When you enter the command, the following information is displayed:

```
Before you can use, extract, or install IBM MQ V9.0, you must accept
the terms of 1. IBM International License Agreement for Evaluation of
Programs 2. IBM International Program License Agreement and additional
license information. Please read the following license agreements carefully.
```

```
The license agreement is separately viewable using the
--viewLicenseAgreement option.
```

```
Press Enter to display the license terms now, or 'x' to skip.
```

3. Review and accept the license terms:
  - a. To display the license, press Enter. Alternatively, pressing x skips the display of the license. After the license is displayed, or immediately if you select x, the following message is displayed:

```
Additional license information is separately viewable using the
--viewLicenseInfo option.
```

```
Press Enter to display additional license information now, or 'x' to skip.
```

- b. To display the additional license terms, press Enter. Alternatively, pressing x skips the display of the additional license terms. After the additional license terms are displayed, or immediately if you select x, the following message is displayed:

```
By choosing the "I Agree" option below, you agree to the terms of the
license agreement and non-IBM terms, if applicable. If you do not
agree, select "I do not Agree".
```

```
Select [1] I Agree, or [2] I do not Agree:
```

- c. To accept the license agreement and continue with selecting the installation directory, select 1. Alternatively, selecting 2 ends the installation immediately. If you select 1, the following message appears:

```
Enter directory for product files or leave blank to accept the default value.
The default target directory is H:\WMQ
```

```
Target directory for product files?
```

4. Specify the installation directory for the resource adapter:
  - If you want to install the product files in the default location, press Enter without specifying a value.
  - If you want to install the product files in a different location from the default, specify the name of the directory in which you want to install the product files, and then press Enter to start the installation.


The directory name that you specify must not already exist, otherwise, when you start the installation, an error is reported and no files are installed. Provided that it does not already exist, the selected installation directory is created and the program files are installed in this directory. During the installation, a new directory with the name `wmq` is created within the installation directory that you selected. Three subdirectories, `JavaEE`, `JavaSE`, and `OSGi`, are created in the `wmq` directory with the following contents:

```
. \JavaEE:
wmq.jmsra.ivt.ear
wmq.jmsra.rar

. \JavaSE:
com.ibm.mq.allclient.jar
com.ibm.mq.traceControl.jar
fscontext.jar
jms.jar
providerutil.jar

. \OSGi:
com.ibm.mq.osgi.allclient_V.R.M.F.jar
com.ibm.mq.osgi.allclientprereqs_V.R.M.F.jar
```

where *V.R.M.F* is the Version, Release, Modification, and Fix Pack number.

 Before IBM MQ Version 9.0.0, Fix Pack 3 and IBM MQ Version 9.0.5, the files that are installed in the `JavaSE` directory include the `JSON4J.jar` file. However, this JAR file is not required, and is therefore removed from the `V.R.M.F-WS-MQ-Install-Java-All.jar` file from Version 9.0.0, Fix Pack 3 and Version 9.0.5. Also, from Version 9.0.0, Fix Pack 3 and Version 9.0.5, there are two changes to the `com.ibm.mq.allclient.jar` file:

- The reference to `JSON4J.jar` file is removed from the class path statement within the manifest file for the `com.ibm.mq.allclient.jar` file.
- The package `com.ibm.msg.client.mqlight` is no longer included inside the `com.ibm.mq.allclient.jar` file.

When the installation is complete, a confirmation message is displayed as shown in the following example:

```
Extracting files to H:\WMQ\wmq
Successfully extracted all product files.
```

## Overview of IBM MQ classes for JMS object pooling

Since the introduction of IBM WebSphere MQ Version 7.0 there has been an increase in some stand-alone applications using frameworks, or being deployed into cloud environments. An increase in server consolidation of applications and queue managers has also seen a greater number of client connections into QueueManagers. A form of connection pooling outside of Java EE helps to reduce overall load.

Within the Java EE programming model, there is a well defined life cycle of the various objects in use. Message-driven beans (MDBs) are most constrained, whilst Servlets provide more freedom. Therefore, the pooling options that are available within the Java EE servers suit the various programming models used.

With Java SE (or with another framework such as Spring) the programming models are extremely flexible. Therefore a single pooling strategy does not suit all. You should consider if there is going to be a framework in place that could do any form of pooling, for example, Spring.

The pooling strategy to use depends on the environment in which your application is running:

“Using IBM MQ classes for JMS” on page 824

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the `javax.jms` package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

### **Object pooling in a Java EE environment:**

Java EE application servers provide connection pooling functionality that can be used by message-driven bean applications, Enterprise Java Beans and Servlets.

For information about how this pooling works in WebSphere Application Server, see [Configuring connection pooling for JMS connections](#).

For more information, see:

- Use JMS connection pools
- Advanced connection pool properties
- Examples of using the connection pool
- Free connection pool maintenance threads
- Pool maintenance thread examples
- JMS connections and IBM MQ

For other application servers, refer to the appropriate application server documentation

“Using IBM MQ classes for JMS” on page 824

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the `javax.jms` package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

### **Object pooling in a Java SE environment:**

With Java SE (or with another framework such as Spring) the programming models are extremely flexible. Therefore a single pooling strategy does not suit all. You should consider if there is going to be a framework in place that could do any form of pooling, for example, Spring.

Otherwise, application logic could take this up. Ask yourself how complex is the application itself? It is best to understand the application and what it demands from the connectivity to the messaging system. Applications are often written as well within their own wrapper code around the basic JMS API.

Whilst this can be a very sensible approach, and can hide complexity, it is worth keeping in mind that it can introduce problems. For example, a generic `getMessage()` method, that is frequently called, should not just open and close consumers.

Points you should consider:

- How long will the application need access to IBM MQ? All the time, or just occasionally.
- How often will messages be sent? The less frequently, the more a single connection to IBM MQ could be shared.
- A connection broken exception is usually a sign of needing to re-create a pooled connection. What about:
  - Security exceptions or host not available
  - Queue full exceptions
- If a connection broken exception occurs, what should happen to the other free connections in the pool? Should they be closed off and re-created?
- If TLS is being used, for example, how long do you want a single connection to remain open?

- How will a pooled connection identify itself such that a queue manager administrator can spot the connection and track it back.

You should consider all JMS objects for pooling, and pool that object whenever it is possible to do so. The objects include:

- JMS connections
- Session
- Contexts
- Producers and consumers of all different types

When using the client transport, JMS connections, sessions, and contexts, will use sockets when communicating with the IBM MQ queue manager. By pooling these objects, the savings are on the number of incoming IBM MQ connections (hConns) to the queue manager and a reduction in the number of channel instances.

Using the bindings transport to the queue manager removes the networking layer entirely. However, many applications use the client transport to provide a more highly available, and workload balanced, configuration.

JMS producers and consumers open destinations on the queue manager. If fewer numbers of queues or topics are opened, and multiple parts of the application are using these objects, pooling these can be useful.

From an IBM MQ perspective, this process saves a sequence of MQOPEN and MQCLOSE operations.

### **Connections, sessions, and contexts**

These objects all encapsulate IBM MQ connection handles to the queue manager, and are generated from a `ConnectionFactory`. You can add logic to an application to constrain the number of connections, and other objects, created from a single connection factory to a specific number.

You can use a simple data structure in the application to contain the connections that are created. The application code that needs to use one of these data structures can *check-out* an object to use.

Take the following factors into account:

- When should connections be removed from the pool? Generally, create an exception listener on the connection. When that listener is called to process an exception, you should re-create the connection, and any sessions created from that connection.
- If a CCDT is in use for workload balancing, the connections could go to different queue managers. This might be applicable for the pooling requirements.

Remember, that the JMS specification states that it is a programming error for multiple threads to be accessing a session or context at the same time. The IBM MQ JMS code does attempt to be rigorous in its handling of threads. However, you should add logic to the application, to ensure that a session or context object is only used by one thread at a time.

### **Producers and consumers**

Each producer and consumer that is created opens a destination on the queue manager. If the same destination is going to be used for a variety of tasks, it makes sense to keep the consumer or producer objects open. Only close the object when all the work is done.

Although opening and closing a destination are short operations, if they are done frequently the time taken can add up.

The scope of these objects is within the session or context they are created from, therefore, they need to be held within that scope. Generally, applications are written such that this is quite straightforward to do.

## Monitoring

How will the applications monitor their object pools? The answer to this is largely determined by the complexity of the solution of pooling implemented.

If you consider a JavaEE pooling implementation, there are a large number of options, including the:

- Current size of the pools
- Time objects have spent in them
- Cleaning of the pools
- Refreshing of the connections

You should also consider how a single re-used session appears on the queue manager. There are connection factory properties to identify the application (such as `appName`) that could be useful.

“Using IBM MQ classes for JMS” on page 824

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the `javax.jms` package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

## Introduction to automatic client reconnection in Java EE environments

The IBM MQ Resource Adapter, which can be deployed into Java EE environments, and the WebSphere Application Server messaging provider use the IBM MQ classes for JMS to communicate with IBM MQ queue managers.

The options available to provide automatic client reconnection in a Java EE environment are:

- Activation specification
- WebSphere Application Server listener ports
- Enterprise JavaBeans and web-based applications
- Applications running inside client containers

For more information see Automatic client reconnection in Java EE environments

Scenarios: WebSphere Application Server with IBM MQ

Scenarios: WebSphere Application Server Liberty profile with IBM MQ

## Character string conversions in IBM MQ classes for JMS

From IBM MQ Version 8.0, some of the default behavior regarding character string conversion in the IBM MQ classes for JMS has changed.

Before IBM MQ Version 8.0, string conversions in IBM MQ classes for JMS was done by calling the `java.nio.charset.Charset.decode(ByteBuffer)` and `Charset.encode(CharBuffer)` methods.

Using either of these methods results in a default replacement ( `REPLACE`) of malformed or untranslatable data.

This behavior can obscure errors in applications, and lead to unexpected characters, for example `?`, in translated data. From IBM MQ Version 8.0, to detect such issues earlier and more effectively, the IBM MQ classes for JMS use `CharsetEncoders` and `CharsetDecoders` directly and configure the handling of malformed and untranslatable data explicitly.

From IBM MQ Version 8.0, the default behavior is to `REPORT` such issues by throwing a suitable `MQException` or `JMSException`.

## Configuring

Translating from UTF-16 (the character representation used in Java) to a native character set, such as UTF-8, is termed encoding, while translating in the opposite direction is termed decoding.

Currently, decoding takes the default behavior for `CharsetDecoders`, reporting errors by throwing an exception.

One setting is used to specify a `java.nio.charset.CodingErrorAction` to control error handling on both encoding and decoding. One other setting is used to control the replacement byte, or bytes, when encoding. The default Java replacement String will be used in decoding operations.

### IBM MQ Classes for JMS

From IBM MQ Version 8.0, two new properties are available. The appropriate constant definitions are in `com.ibm.msg.client.wmq.WMQConstants`

#### JMS\_IBM\_UNMAPPABLE\_ACTION

Sets or gets the `CodingErrorAction` to apply when a character cannot be mapped in an encoding or decoding operation.

You should set this as `CodingErrorAction.{REPLACE|REPORT|IGNORE}.toString()` as follows:

```
public static final String JMS_IBM_UNMAPPABLE_ACTION = "JMS_IBM_Unmappable_Action";
```

#### JMS\_IBM\_UNMAPPABLE\_REPLACEMENT

Sets or gets the replacement bytes to apply when a character cannot be mapped in an encoding operation.

The default Java replacement String is used in decoding operations.

```
public static final String JMS_IBM_UNMAPPABLE_REPLACEMENT = "JMS_IBM_Unmappable_Replacement";
```

The `JMS_IBM_UNMAPPABLE_ACTION` and `JMS_IBM_UNMAPPABLE_REPLACEMENT` properties can be set on destinations or messages. A value set on a message overrides the value set on the destination to which the message is being sent.

Note that `JMS_IBM_UNMAPPABLE_REPLACEMENT` must be set as a single byte.

### Setting system defaults

From IBM MQ Version 8.0, the following two Java system properties are available to configure default behavior regarding character string conversion. Note that destinations or messages override these properties.

#### `com.ibm.mq.cfg.jmqi.UnmappableCharacterAction`

Specifies the action to be taken for untranslatable data on encoding and decoding. The value can be `REPORT`, `REPLACE`, or `IGNORE`.

#### `com.ibm.mq.cfg.jmqi.UnmappableCharacterReplacement`

Sets or gets the replacement bytes to apply when a character cannot be mapped in an encoding operation. The default Java replacement string is used in decoding operations.

To avoid confusion between Java character and native byte representations, you should specify `com.ibm.mq.cfg.jmqi.UnmappableCharacterReplacement` as a decimal number representing the replacement byte in the native character set.

For example, the decimal value of `?`, as a native byte, is 63 if the native character set is ASCII-based, such as ISO-8859-1, while it is 111 if the native character set is EBCDIC.

## Related concepts:

“Character string conversions in IBM MQ classes for Java” on page 1025

From IBM MQ Version 8.0, some of the default behavior regarding character string conversion in the IBM MQ classes for Java has changed.

## Writing IBM MQ classes for JMS applications

After a brief introduction to the JMS model, this topic provides detailed guidance on how to write IBM MQ classes for JMS applications.

### The JMS model:

The JMS model defines a set of interfaces that Java applications can use to perform messaging operations. IBM MQ classes for JMS, as a JMS provider, defines how JMS objects are related to IBM MQ concepts. The JMS specification expects certain JMS objects to be administered objects. JMS 2.0 introduces a simplified API, while also retaining the classic API, from JMS 1.1.

The JMS specification and the `javax.jms` package define a set of interfaces that Java applications can use to perform messaging operations.

From Version 8.0, IBM MQ supports the JMS 2.0 version of the JMS standard, which introduces a simplified API, while also retaining the classic API, from JMS 1.1.

### Simplified API

JMS 2.0 introduces the simplified API, while also retaining the domain specific and domain independent interfaces from JMS 1.1. The simplified API reduces the number of objects that are needed to send and receive messages and consists of the following interfaces:

#### ConnectionFactory

A `ConnectionFactory` is an administered object that is used by a JMS client to create a `Connection`. This interface is also used in the classic API.

#### JMSContext

This object combines the `Connection` and `Session` objects of the classic API. `JMSContext` objects can be created from other `JMSContext` objects, with the underlying connection being duplicated.

#### JMSProducer

A `JMSProducer` is created by a `JMSContext` and is used to send messages to a queue or topic. The `JMSProducer` object causes the creation of objects that are required to send the message.

#### JMSConsumer

A `JMSConsumer` is created by a `JMSContext` and is used to receive messages from a topic or a queue.

The simplified API has a number of effects:

- The `JMSContext` object always automatically starts the underlying connection.
- `JMSProducers` and `JMSConsumers` can now work directly with message bodies, without having to get the whole message object, by using the `Message`'s `getBody` method.
- Message properties can be set on the `JMSProducer` object, using method chaining, before sending a 'body', a messages content. The `JMSProducer` will handle the creation of all objects that are needed to send the message. Using JMS 2.0, properties can be set, and a message sent as follows:

```
context.createProducer().
setProperty("foo", "bar").
setTimeToLive(10000).
setDeliveryMode(NON_PERSISTENT).
setDisableMessageTimestamp(true).
send(dataQueue, body);
```

JMS 2.0 also introduces shared subscriptions where messages can be shared between multiple consumers. All JMS 1.1 subscriptions are treated as unshared subscriptions.

## **Classic API**

The following list summarizes the main JMS interfaces of the classic API:

### **Destination**

A destination is where an application sends messages, or it is a source from which an application receives messages, or both.

### **ConnectionFactory**

A ConnectionFactory object encapsulates a set of configuration properties for a connection. An application uses a connection factory to create a connection.

### **Connection**

A Connection object encapsulates an application's active connection to a messaging server. An application uses a connection to create sessions.

### **Session**

A session is a single threaded context for sending and receiving messages. An application uses a session to create messages, message producers, and message consumers. A session is either transacted or not transacted.

### **Message**

A Message object encapsulates a message that an application sends or receives.

### **MessageProducer**

An application uses a message producer to send messages to a destination.

### **MessageConsumer**

An application uses a message consumer to receive messages sent to a destination.

Figure 107 on page 873 shows these objects and their relationships.



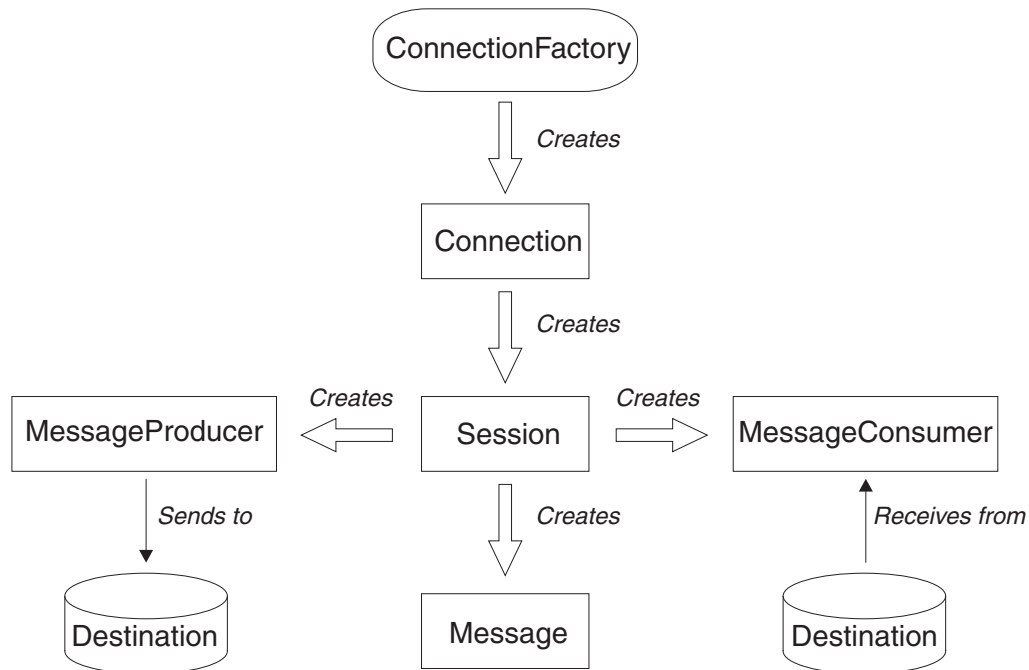


Figure 107. JMS objects and their relationships

A Destination, ConnectionFactory, or Connection object can be used concurrently by different threads of a multithreaded application, but a Session, MessageProducer, or MessageConsumer object cannot be used concurrently by different threads. The simplest way of ensuring that a Session, MessageProducer, or MessageConsumer object is not used concurrently is to create a separate Session object for each thread.

JMS support two styles of messaging:

- Point-to-point messaging
- Publish/subscribe messaging

These styles of messaging are also referred to as *messaging domains*, and you can combine both styles of messaging in an application. In the point-to-point domain, a destination is a queue and, in the publish/subscribe domain, a destination is a topic.

With versions of JMS before JMS 1.1, programming for the point-to-point domain uses one set of interfaces and methods, and programming for the publish/subscribe domain uses another set. The two sets are similar, but separate. From JMS 1.1, you can use a common set of interfaces and methods that support both messaging domains. The common interfaces provide a domain independent view of each messaging domain. Table 113 lists the JMS domain independent interfaces and their corresponding domain specific interfaces.

Table 113. The JMS domain independent and domain specific interfaces

Domain independent interfaces	Domain specific interfaces for the point-to-point domain	Domain specific interfaces for the publish/subscribe domain
ConnectionFactory	QueueConnectionFactory	TopicConnectionFactory
Connection	QueueConnection	TopicConnection
Destination	Queue	Topic
Session	QueueSession	TopicSession
MessageProducer	QueueSender	TopicPublisher

Table 113. The JMS domain independent and domain specific interfaces (continued)

Domain independent interfaces	Domain specific interfaces for the point-to-point domain	Domain specific interfaces for the publish/subscribe domain
MessageConsumer	QueueReceiver QueueBrowser	TopicSubscriber

JMS 2.0 retains all the domain specific interfaces, and so existing applications can still use these interfaces. For new applications, however, consider using the domain independent interfaces of JMS 1.1 or the simplified API of JMS 2.0.

In IBM MQ classes for JMS, JMS objects are related to IBM MQ concepts in the following ways:

- A Connection object has properties that are derived from the properties of the connection factory that was used to create the connection. These properties control how an application connects to a queue manager. Examples of these properties are the name of the queue manager and, for an application that connects to the queue manager in client mode, the host name or IP address of the system on which the queue manager is running.
- A Session object encapsulates an IBM MQ connection handle, which therefore defines the transactional scope of the session.
- A MessageProducer object and a MessageConsumer object each encapsulates an IBM MQ object handle.

When using IBM MQ classes for JMS, all the normal rules of IBM MQ apply. Note, in particular, that an application can send a message to a remote queue but it can receive a message only from a queue that is owned by the queue manager to which the application is connected.

The JMS specification expects ConnectionFactory and Destination objects to be administered objects. An administrator creates and maintains administered objects in a central repository, and a JMS application retrieves these objects using the Java Naming and Directory Interface (JNDI).

In IBM MQ classes for JMS, the implementation of the Destination interface is an abstract superclass of Queue and Topic, and so an instance of Destination is either a Queue object or a Topic object. The domain independent interfaces treat a queue or a topic as a destination. The messaging domain for a MessageProducer or MessageConsumer object is determined by whether the destination is a queue or a topic.

In IBM MQ classes for JMS therefore, objects of the following types can be administered objects:

- ConnectionFactory
- QueueConnectionFactory
- TopicConnectionFactory
- Queue
- Topic
- XAConnectionFactory
- XAQueueConnectionFactory
- XATopicConnectionFactory

**Related concepts:**

“Using JMS 2.0 functionality” on page 996

JMS 2.0 introduces several new areas of functionality to the IBM MQ classes for JMS.

**Related information:**

IBM MQ Java language interfaces

**JMS messages:**

JMS messages are composed of a header, properties, and a body. JMS defines five types of message body.

JMS messages are composed of the following parts:

**Header**

All messages support the same set of header fields. Header fields contain values that are used by both clients and providers to identify and route messages.

**Properties**

Each message contains a built-in facility to support application-defined property values. Properties provide an efficient mechanism to filter application-defined messages.

**Body** JMS defines five types of message body that cover the majority of messaging styles currently in use:

**Stream**

A stream of Java primitive values. It is filled and read sequentially.

**Map**

A set of name-value pairs, where names are strings and values are Java primitive types. The entries can be accessed sequentially or randomly by name. The order of the entries is undefined.

**Text**

A message containing a `java.lang.String`.

**Object**

A message that contains a serializable Java object

**Bytes**

A stream of uninterpreted bytes. This message type is for literally encoding a body to match an existing message format.

The `JMSCorrelationID` header field is used to link one message with another. It typically links a reply message with its requesting message. `JMSCorrelationID` can hold a provider-specific message ID, an application-specific String, or a provider-native `byte[]` value.

*Message selectors in JMS:*

Messages can contain application-defined property values. An application can use message selectors to have a JMS provider filter messages.

A message contains a built-in facility to support application-defined property values. In effect, this provides a mechanism to add application-specific header fields to a message. Properties allow an application, using message selectors, to have a JMS provider select or filter messages on its behalf, using application-specific criteria. Application-defined properties must obey the following rules:

- Property names must obey the rules for a message selector identifier.
- Property values can be Boolean, byte, short, int, long, float, double, and String.
- The `JMSX` and `JMS_` name prefixes are reserved.

Property values are set before sending a message. When a client receives a message, the message properties are read-only. If a client attempts to set properties at this point, a `MessageNotWriteableException` is thrown. If `clearProperties` is called, the properties can now be both read from, and written to.

A property value might duplicate a value in a message body. JMS does not define a policy for what might be made into a property. However, application developers must be aware that JMS providers probably handle data in a message body more efficiently than data in message properties. For best performance, applications must use message properties only when they need to customize a message header. The primary reason for doing this is to support customized message selection.

A JMS message selector allows a client to specify the messages that it is interested in by using the message header. Only messages with headers that match the selector are delivered.

Message selectors cannot refer to message body values.

A message selector matches a message when the selector evaluates to true when the message header field and property values are substituted for their corresponding identifiers in the selector.

A message selector is a String, with syntax that is based on a subset of the SQL92 conditional expression syntax. The order in which a message selector is evaluated is from left to right within a precedence level. You can use parentheses to change this order. Predefined selector literals and operator names are written here in uppercase; however, they are not case-sensitive.

### Contents of a message selector

A message selector can contain:

- Literals
  - A string literal is enclosed in quotation marks. A doubled quotation mark represents a quotation mark. Examples are 'literal' and 'literal"s'. Like Java string literals, these use the Unicode character encoding.
  - An exact numeric literal is a numeric value without a decimal point, such as 57, -957, and +62. Numbers in the range of Java long are supported.
  - An approximate numeric literal is a numeric value in scientific notation, such as 7E3 or -57.9E2, or a numeric value with a decimal, such as 7., -95.7, or +6.2. Numbers in the range of Java double are supported.
  - The Boolean literals TRUE and FALSE.
- Identifiers:
  - An identifier is an unlimited length sequence of Java letters and Java digits, the first of which must be a Java letter. A letter is any character for which the method `Character.isJavaLetter` returns true. This includes `_` and `$`. A letter or digit is any character for which the method `Character.isJavaLetterOrDigit` returns true.
  - Identifiers cannot be the names NULL, TRUE, or FALSE.
  - Identifiers cannot be NOT, AND, OR, BETWEEN, LIKE, IN, or IS.
  - Identifiers are either header field references or property references.
  - Identifiers are case sensitive.
  - Message header field references are restricted to:
    - JMSDeliveryMode
    - JMSPriority
    - JMSMessageID
    - JMSTimestamp
    - JMSCorrelationID
    - JMSTypeJMSMessageID, JMSTimestamp, JMSCorrelationID, and JMSType values can be null, and if so, are treated as a NULL value.
  - Any name beginning with JMSX is a JMS-defined property name.

- Any name beginning with JMS\_ is a provider-specific property name.
- Any name that does not begin with JMS is an application-specific property name. If there is a reference to a property that does not exist in a message, its value is NULL. If it does exist, its value is the corresponding property value.
- White space is the same as it is defined for Java: space, horizontal tab, form feed, and line terminator.
- Expressions:
  - A selector is a conditional expression. A selector that evaluates to true matches; a selector that evaluates to false or unknown does not match.
  - Arithmetic expressions are composed of themselves, arithmetic operations, identifiers (with a value that is treated as a numeric literal), and numeric literals.
  - Conditional expressions are composed of themselves, comparison operations, and logical operations.
- Standard bracketing (), to set the order in which expressions are evaluated, is supported.
- Logical operators in precedence order: NOT, AND, OR.
- Comparison operators: =, >, >=, <, <=, <> (not equal).
  - Only values of the same type can be compared. One exception is that it is valid to compare exact numeric values and approximate numeric values. (The type conversion required is defined by the rules of Java numeric promotion.) If there is an attempt to compare different types, the selector is always false.
  - String and Boolean comparison is restricted to = and <>. Two strings are equal only if they contain the same sequence of characters.
- Arithmetic operators in precedence order:
  - +, - unary.
  - \*, /, multiplication, and division.
  - +, -, addition, and subtraction.
  - Arithmetic operations on a NULL value are not supported. If they are attempted, the complete selector is always false.
  - Arithmetic operations must use Java numeric promotion.
- arithmetic-expr1 [NOT] BETWEEN arithmetic-expr2 and arithmetic-expr3 comparison operator:
  - Age BETWEEN 15 and 19 is equivalent to age >= 15 AND age <= 19.
  - Age NOT BETWEEN 15 and 19 is equivalent to age < 15 OR age > 19.
  - If any of the expressions of a BETWEEN operation are NULL, the value of the operation is false. If any of the expressions of a NOT BETWEEN operation are NULL, the value of the operation is true.
- identifier [NOT] IN (string-literal1, string-literal2,...) comparison operator where identifier has a String or NULL value.
  - Country IN ('UK', 'US', 'France') is true for 'UK' and false for 'Peru'. It is equivalent to the expression (Country = 'UK') OR (Country = 'US') OR (Country = 'France').
  - Country NOT IN ('UK', 'US', 'France') is false for 'UK' and true for 'Peru'. It is equivalent to the expression NOT ((Country = 'UK') OR (Country = 'US') OR (Country = 'France')).
  - If the identifier of an IN or NOT IN operation is NULL, the value of the operation is unknown.
- identifier [NOT] LIKE pattern-value [ESCAPE escape-character] comparison operator, where identifier has a string value. pattern-value is a string literal, where \_ stands for any single character and % stands for any sequence of characters (including the empty sequence). All other characters stand for themselves. The optional escape-character is a single character string literal, with a character that is used to escape the special meaning of the \_ and % in pattern-value.
  - phone LIKE '12%3' is true for 123 and 12993 and false for 1234.
  - word LIKE 'l\_se' is true for "lose" and false for "loose".
  - underscored LIKE '\\_%' ESCAPE '\' is true for "\_foo" and false for "bar".
  - phone NOT LIKE '12%3' is false for 123 and 12993 and true for 1234.



### Mapping JMS messages onto IBM MQ messages:

IBM MQ messages are composed of a Message Descriptor, an optional MQRFH2 header, and a body. The contents of a JMS message are partly mapped and partly copied to an IBM MQ message.

This topic describes how the JMS message structure that is described in the first part of this section is mapped onto an IBM MQ message. It is of interest to programmers who want to transmit messages between JMS and traditional IBM MQ applications. It is also of interest to people who want to manipulate messages transmitted between two JMS applications, for example, in an IBM Integration Bus implementation.

This section does not apply if an application uses a real-time connection to a broker. When an application uses a real-time connection, all communication is performed directly over TCP/IP; no IBM MQ queues or messages are involved.

IBM MQ messages are composed of three components:

- The IBM MQ Message Descriptor (MQMD)
- An IBM MQ MQRFH2 header
- The message body.

The MQRFH2 is optional, and its inclusion in an outgoing message is governed by the TARGCLIENT flag in the JMS Destination class. You can set this flag using the IBM MQ JMS administration tool. Because the MQRFH2 carries JMS-specific information, always include it in the message when the sender knows that the receiving destination is a JMS application. Normally, omit the MQRFH2 when sending a message directly to a non-JMS application. This is because such an application does not expect an MQRFH2 in its IBM MQ message.

If an incoming message does not have an MQRFH2 header, the Queue or Topic object derived from the JMSReplyTo header field of the message, by default, has this flag set so that a reply message sent to the queue or topic also does not have an MQRFH2 header. You can switch off this behavior of including an MQRFH2 header in a reply message only if the original message has an MQRFH2 header, by setting the TARGCLIENTMATCHING property of the connection factory to NO.

Figure 108 shows how the structure of a JMS message is transformed to an IBM MQ message and back again:

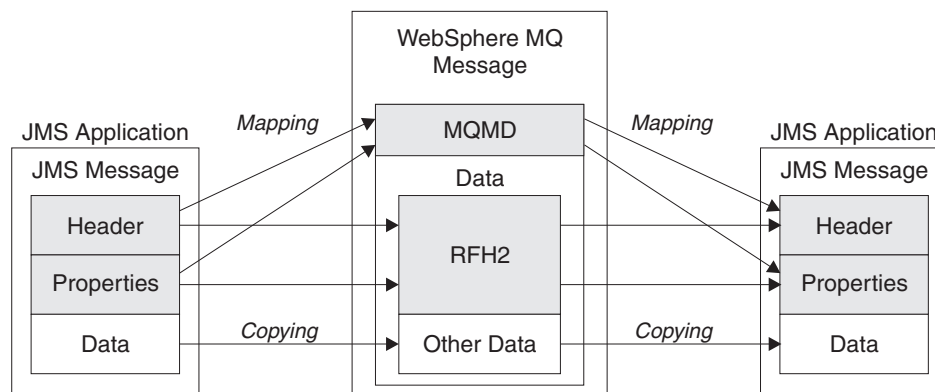


Figure 108. How messages are transformed between JMS and IBM MQ using the MQRFH2 header

The structures are transformed in two ways:

#### Mapping

Where the MQMD includes a field that is equivalent to the JMS field, the JMS field is mapped

onto the MQMD field. Additional MQMD fields are exposed as JMS properties, because a JMS application might need to get or set these fields when communicating with a non-JMS application.

### Copying

Where there is no MQMD equivalent, a JMS header field or property is passed, possibly transformed, as a field inside the MQRFH2.

*The MQRFH2 header and JMS:*

This collection of topics describes the MQRFH Version 2 header, which carries JMS-specific data that is associated with the message content. The MQRFH Version 2 header is extensible, and can also carry additional information that is not directly associated with JMS. However, this section covers only its use by JMS.

For a full description see MQRFH2 - Rules and formatting header 2.

There are two parts of the header, a fixed portion and a variable portion.

### Fixed portion

The fixed portion is modeled on the *standard* IBM MQ header pattern and consists of the following fields:

#### StrucId (MQCHAR4)

Structure identifier.

Must be MQRFH\_STRUC\_ID (value: "RFH ") (initial value).

MQRFH\_STRUC\_ID\_ARRAY (value: "R", "F", "H", " ") is also defined.

#### Version (MQLONG)

Structure version number.

Must be MQRFH\_VERSION\_2 (value: 2) (initial value).

#### StrucLength (MQLONG)

Total length of MQRFH2, including the NameValueData fields.

The value set into StrucLength must be a multiple of 4 (the data in the NameValueData fields can be padded with space characters to achieve this).

#### Encoding (MQLONG)

Data encoding.

Encoding of any numeric data in the portion of the message following the MQRFH2 (the next header, or the message data following this header).

#### CodedCharSetId (MQLONG)

Coded character set identifier.

Representation of any character data in the portion of the message following the MQRFH2 (the next header, or the message data following this header).

#### Format (MQCHAR8)

Format name.

Format name for the portion of the message following the MQRFH2.

#### Flags (MQLONG)

Flags.

MQRFH\_NO\_FLAGS =0. No flags set.

#### NameValueCCSID (MQLONG)

The coded character set identifier (CCSID) for the NameValueData character strings contained in



this header. The NameValueData can be coded in a character set that differs from the other character strings that are contained in the header (StrucID and Format).

If the NameValueCCSID is a 2 byte Unicode CCSID (1200, 13488, or 17584), the byte order of the Unicode is the same as the byte ordering of the numeric fields in the MQRFH2. (For example, Version, StrucLength, and NameValueCCSID itself.)


The NameValueCCSID takes values from the following table: 

Table 114. Possible values for NameValueCCSID field

CCSID	Meaning
1200	UTF-16, most recent Unicode version supported
13488	UTF-16, Unicode version 2.0 subset
17584	UTF-16, Unicode version 3.0 subset (includes the Euro symbol)
1208	UTF-8, most recent Unicode version supported

### Variable portion

The variable portion follows the fixed portion. The variable portion contains a variable number of MQRFH2 folders. Each folder contains a variable number of elements or properties. Folders group related properties. The MQRFH2 headers created by JMS can contain any of the following folders:

#### The mcd folder

mcd contains properties that describe the format of the message. For example, the message service domain Msd property identifies a JMS message as being JMSTextMessage, JMSBytesMessage, JMSStreamMessage, JMSMapMessage, JMSObjectMessage, or null.

The mcd folder is always present in a JMS message containing an MQRFH2.

It is always present in a message containing an MQRFH2 sent from IBM Integration Bus. It describes the domain, format, type, and message set of a message.

Table 115. mcd property name, synonym, data type, and folder

Property synonym	Property name	Data type	Folder
	mcd.Msd	string	<mcd><Msd>messageDomain</Msd></mcd>
	mcd.Set	string	<mcd><Set>messageDomain</Set></mcd>
	mcd.Type	string	<mcd><Type>messageDomain</Type></mcd>
	mcd.Fmt	string	<mcd><Fmt>messageDomain</Fmt></mcd>

Do not add your own properties in the mcd folder.

#### The jms folder

jms contains JMS header fields, and JMSX properties that cannot be fully expressed in the MQMD. The jms folder is always present in a JMS MQRFH2.

#### The usr folder

usr contains application-defined JMS properties associated with the message. The usr folder is present only if an application has set an application-defined property.

#### The mqext folder

mqext contains the following types of property:

- Properties that are used only by WebSphere Application Server.

- Properties relating to delayed delivery of messages.

The folder is present if the application has either set at least one of the IBM defined properties or used delivery delay.

Table 116. *mqext* property name, synonym, data type, and folder

Property synonym	Property name	Data type	Folder
JMSArmCorrelator	mqext.Arm	string	<mqext><Arm>armCorrelator</Arm></mqext>
JMSRMCorrelator	mqext.Wrm	string	<mqext><Wrm>wrmCorrelator</Wrm></mqext>
JMSDeliveryTime	mqext.Dlt	i8	<mqext><Dlt>DeliveryTime</Dlt></mqext>
JMSDeliveryDelay	mqext.Dly	i8	<mqext><Dly>DeliveryTime</Dly></mqext>

Do not add your own properties in the *mqext* folder.

### The *mqs* folder

*mqs* contains properties that are used only by IBM MQ publish/subscribe. The folder is present only if the application has set at least one of the integrated publish/subscribe properties.

Table 117. *mqs* property name, synonym, data type, and folder

Property synonym	Property name	Data type	Folder
MQTopicString	mqs.Top	string	<mqs><Top>topicString</Top></mqs>
MQSubUserData	mqs.Sud	string	<mqs><Sud>subscriberUserData</Sud></mqs>
MQIsRetained	mqs.Ret	boolean	<mqs><Ret>isRetained</Ret></mqs>
MQPubOptions	mqs.Pub	i8	<mqs><Pub>publicationOptions</Pub></mqs>
MQPubLevel	mqs.Pbl	i8	<mqs><Pbl>publicationLevel</Pbl></mqs>
MQPubTime	mqs.Pts	string	<mqs><Pts>publicationTime</Pts></mqs>
MQPubSeqNum	mqs.Seq	i8	<mqs><Seq>publicationSequenceNumber</Seq></mqs>
MQPubStrIntData	mqs.Sid	string	<mqs><Sid>publicationData</Sid></mqs>
MQPubFormat	mqs.Pfmt	i8	<mqs><Pfmt>messageFormat</Pfmt></mqs>

Do not add your own properties in the *mqs* folder.

Table 118 shows a full list of property names.

Table 118. *MQRFH2* folders and properties used by JMS

JMS field name	Java type	MQRFH2 folder name	Property name	Type/values
JMSDestination	Destination	jms	Dst	string
JMSExpiration	long	jms	Exp	i8
JMSPriority	int	jms	Pri	i4
JMSDeliveryMode	int	jms	Dlv	i4
JMSCorrelationID	String	jms	Cid	string
JMSReplyTo	Destination	jms	Rto	string
JMSTimestamp	long	jms	Tms	i8
JMSType	String	mcd	Type, Set, Fmt	string
JMSXGroupID	String	jms	Gid	string
JMSXGroupSeq	int	jms	Seq	i4

Table 118. MQRFH2 folders and properties used by JMS (continued)

JMS field name	Java type	MQRFH2 folder name	Property name	Type/values
xxx (user defined)	Any	usr	xxx	any
		mcd	Msd	jms_none jms_text jms_bytes jms_map jms_stream jms_object

### NameValueLength (MQLONG)

Length in bytes of the NameValueData string that immediately follows this length field (it does not include its own length).

### NameValueData (MQCHARn)

A single character string, whose length in bytes is given by the preceding NameValueLength field. It contains a folder holding a sequence of properties. Each property is a name/type/value triplet, contained within an XML element whose name is the folder name, as follows:

```
<foldername>
triplet1 triplet2 tripletn </foldername>
```

The closing </foldername> tag can be followed by spaces as padding characters. Each triplet is encoded using an XML-like syntax:

```
<name dt='datatype'>value</name>
```

The dt='datatype' element is optional and is omitted for many properties, because the data type is predefined. If it is included, one or more space characters must be included before the dt= tag.

**name** is the name of the property; see Table 118 on page 882.

#### datatype

must match, after folding, one of the data types listed in Table 119.

**value** is a string representation of the value to be conveyed, using the definitions in Table 119.

A null value is encoded using the following syntax:

```
<name dt='datatype' xsi:nil='true'></name>
```

Do not use xsi:nil='false'.

Table 119. Property data types

Data type	Definition
string	Any sequence of characters excluding < and &
boolean	The character 0 or 1 ( 0 = false, 1 = true)
bin.hex	Hexadecimal digits representing octets
i1	A number, expressed using digits 0..9, with optional sign (no fractions or exponent). Must lie in the range -128 to 127 inclusive
i2	A number, expressed using digits 0..9, with optional sign (no fractions or exponent). Must lie in the range -32768 to 32767 inclusive
i4	A number, expressed using digits 0..9, with optional sign (no fractions or exponent). Must lie in the range -2147483648 to 2147483647 inclusive
i8	A number, expressed using digits 0..9, with optional sign (no fractions or exponent). Must lie in the range -9223372036854775808 to 92233720368547750807 inclusive

Table 119. Property data types (continued)

Data type	Definition
int	A number, expressed using digits 0..9, with optional sign (no fractions or exponent). Must lie in the same range as i8. This can be used in place of one of the i* types if the sender does not want to associate a particular precision with the property
r4	Floating point number, magnitude <= 3.40282347E+38,>= 1.175E-37 expressed using digits 0..9, optional sign, optional fractional digits, optional exponent
r8	Floating point number, magnitude <= 1.7976931348623E+308,>= 2.225E-307 expressed using digits 0..9, optional sign, optional fractional digits, optional exponent

A string value can contain spaces. You must use the following escape sequences in a string value:

- &amp; for the & character
- &lt; for the < character

You can use the following escape sequences, but they are not required:

- &gt; for the > character
- &apos; for the ' character
- &quot; for the " character

JMS fields and properties with corresponding MQMD fields:

These tables show the MQMD fields equivalent to JMS header fields, JMS properties, and JMS provider-specific properties.

Table 120 lists the JMS header fields and Table 121 lists the JMS properties that are mapped directly to MQMD fields. Table 122 on page 885 lists the provider-specific properties and the MQMD fields that they are mapped to.

Table 120. JMS header fields mapping to MQMD fields

JMS header field	Java type	MQMD field	C type
JMSDeliveryMode	int	Persistence	MQLONG
JMSExpiration	long	Expiry	MQLONG
JMSPriority	int	Priority	MQLONG
JMSMessageID	String	MsgID	MQBYTE24
JMSTimestamp	long	PutDate PutTime	MQCHAR8 MQCHAR8
JMSCorrelationID	String	CorrelId	MQBYTE24

Table 121. JMS properties mapping to MQMD fields

JMS property	Java type	MQMD field	C type
JMSXUserID	String	UserIdentifier	MQCHAR12
JMSXAppID	String	PutApplName	MQCHAR28
JMSXDeliveryCount	int	BackoutCount	MQLONG
JMSXGroupID	String	GroupId	MQBYTE24
JMSXGroupSeq	int	MsgSeqNumber	MQLONG

Table 122. JMS provider-specific properties mapping to MQMD fields

JMS provider-specific property	Java type	MQMD field	C type
JMS_IBM_Report_Exception	int	Report	MQLONG
JMS_IBM_Report_Expiration	int	Report	MQLONG
JMS_IBM_Report_COA	int	Report	MQLONG
JMS_IBM_Report_COD	int	Report	MQLONG
JMS_IBM_Report_PAN	int	Report	MQLONG
JMS_IBM_Report_NAN	int	Report	MQLONG
JMS_IBM_Report_Pass_Msg_ID	int	Report	MQLONG
JMS_IBM_Report_Pass_Correl_ID	int	Report	MQLONG
JMS_IBM_Report_Discard_Msg	int	Report	MQLONG
JMS_IBM_MsgType	int	MsgType	MQLONG
JMS_IBM_Feedback	int	Feedback	MQLONG
JMS_IBM_Format	String	Format <sup>1</sup>	MQCHAR8
JMS_IBM_PutApplType	int	PutApplType	MQLONG
JMS_IBM_Encoding	int	Encoding	MQLONG
JMS_IBM_Character_Set	String	CodedCharacterSetId <sup>2</sup>	MQLONG
JMS_IBM_PutDate	String	PutDate	MQCHAR8
JMS_IBM_PutTime	String	PutTime	MQCHAR8
JMS_IBM_Last_Msg_In_Group	boolean	MsgFlags	MQLONG

**Note:**

1. JMS\_IBM\_Format represents the format of the message body. This can be defined by the application setting the JMS\_IBM\_Format property of the message (note that there is an 8 character limit), or can default to the IBM MQ format of the message body appropriate to the JMS message type. JMS\_IBM\_Format maps to the MQMD Format field only if the message contains no RFH or RFH2 sections. In a typical message, it maps to the Format field of the RFH2 immediately preceding the message body.
2. JMS\_IBM\_Character\_Set property value is a String value that contains the Java character set equivalent for the numeric CodedCharacterSetId value. MQMD field CodedCharacterSetId is a numeric value that contains the equivalent of the Java character set string specified by the JMS\_IBM\_Character\_Set property.

*Mapping JMS fields onto IBM MQ fields (outgoing messages):*

These tables show how JMS header and property fields are mapped into MQMD and MQRFH2 fields at send() or publish() time.

Table 123 shows how the JMS header fields are mapped into MQMD/RFH2 fields at send() or publish() time. Table 124 shows how JMS properties are mapped into MQMD/RFH2 fields at send() or publish() time. Table 125 on page 887 shows how JMS provider-specific properties are mapped to MQMD fields at send() or publish() time,

For fields marked Set by Message Object, the value transmitted is the value held in the JMS message immediately before the send() or publish() operation. The value in the JMS message is left unchanged by the operation.

For fields marked Set by Send Method, a value is assigned when the send() or publish() is performed (any value held in the JMS message is ignored). The value in the JMS message is updated to show the value used.

Fields marked as Receive-only are not transmitted and are left unchanged in the message by send() or publish().

*Table 123. Outgoing message field mapping*

JMS header field name	MQMD field used for transmission	Header	Set by
JMSDestination		MQRFH2	Send Method
JMSDeliveryMode	Persistence	MQRFH2	Send Method
JMSExpiration	Expiry	MQRFH2	Send Method
JMSPriority	Priority	MQRFH2	Send Method
JMSMessageID	MsgID		Send Method
JMSTimestamp	PutDate/PutTime		Send Method
JMSCorrelationID	CorrelId	MQRFH2	Message Object
JMSReplyTo	ReplyToQ/ReplyToQMgr	MQRFH2	Message Object
JMSType		MQRFH2	Message Object
JMSRedelivered			Receive-only

**Note:**

1. MQMD field CodedCharacterSetId is a numeric value that contains the equivalent of the Java character set string specified by the JMS\_IBM\_Character\_Set property.

*Table 124. Outgoing message JMS property mapping*

JMS property name	MQMD field used for transmission	Header	Set by
JMSXUserID	UserIdentifier		Send Method
JMSXAppID	PutApplName		Send Method
JMSXDeliveryCount			Receive-only
JMSXGroupID	GroupId	MQRFH2	Message Object
JMSXGroupSeq	MsgSeqNumber	MQRFH2	Message Object

Table 125. Outgoing message JMS provider-specific property mapping

JMS provider-specific property name	MQMD field used for transmission	Header	Set by
JMS_IBM_Report_Exception	Report		Message Object
JMS_IBM_Report_Expiration	Report		Message Object
JMS_IBM_Report_COA/COD	Report		Message Object
JMS_IBM_Report_NAN/PAN	Report		Message Object
JMS_IBM_Report_Pass_Msg_ID	Report		Message Object
JMS_IBM_Report_Pass_Correl_ID	Report		Message Object
JMS_IBM_Report_Discard_Msg	Report		Message Object
JMS_IBM_MsgType	MsgType		Message Object
JMS_IBM_Feedback	Feedback		Message Object
JMS_IBM_Format	Format		Message Object
JMS_IBM_PutApplType	PutApplType		Send Method
JMS_IBM_Encoding	Encoding		Message Object
JMS_IBM_Character_Set	CodedCharacterSetId		Message Object
JMS_IBM_PutDate	PutDate		Send Method
JMS_IBM_PutTime	PutTime		Send Method
JMS_IBM_Last_Msg_In_Group	MsgFlags		Message Object

Mapping JMS header fields at *send()* or *publish()*:

These notes relate to the mapping of JMS fields at *send()* or *publish()*.

#### JMSDestination to MQRFH2

This is stored as a string that serializes the salient characteristics of the destination object, so that a receiving JMS can reconstitute an equivalent destination object. The MQRFH2 field is encoded as URI (see “Uniform resource identifiers (URIs)” on page 946 for details of the URI notation).

#### JMSReplyTo to MQMD.ReplyToQ, ReplyToQMgr, MQRFH2

The queue name is copied to the MQMD.ReplyToQ field, and the queue manager name is copied to the ReplyToQMgr fields. The destination extension information (other useful details that are kept in the destination object) is copied into the MQRFH2 field. The MQRFH2 field is encoded as a URI (see “Uniform resource identifiers (URIs)” on page 946 for details of the URI notation).

#### JMSDeliveryMode to MQMD.Persistence

The JMSDeliveryMode value is set by the *send()* or *publish()* Method or MessageProducer, unless the Destination Object overrides it. The JMSDeliveryMode value is mapped to the MQMD.Persistence field as follows:

- JMS value PERSISTENT is equivalent to MQPER\_PERSISTENT
- JMS value NON\_PERSISTENT is equivalent to MQPER\_NOT\_PERSISTENT

If the MQQueue persistence property is not set to WMQConstants.WMQ\_PER\_QDEF, the delivery mode value is also encoded in the MQRFH2.

#### JMSExpiration to/from MQMD.Expiry, MQRFH2

JMSExpiration stores the time to expire (the sum of the current time and the time to live), whereas MQMD stores the time to live. Also, JMSExpiration is in milliseconds, but MQMD.Expiry is in tenths of a second.

- If the *send()* method sets an unlimited time to live, MQMD.Expiry is set to MQEI\_UNLIMITED, and no JMSExpiration is encoded in the MQRFH2.

- If the `send()` method sets a time to live that is less than 214748364.7 seconds (about 7 years), the time to live is stored in `MQMD.Expiry`, and the expiration time (in milliseconds), is encoded as an i8 value in the `MQRFH2`.
- If the `send()` method sets a time to live greater than 214748364.7 seconds, `MQMD.Expiry` is set to `MQEI_UNLIMITED`. The true expiration time in milliseconds is encoded as an i8 value in the `MQRFH2`.

### **JMSPriority to MQMD.Priority**

Directly map `JMSPriority` value (0-9) onto `MQMD` priority value (0-9). If `JMSPriority` is set to a non-default value, the priority level is also encoded in the `MQRFH2`.

### **JMSMessageID from MQMD.MessageID**

All messages sent from JMS have unique message identifiers assigned by IBM MQ. The value assigned is returned in the `MQMD.MessageId` field after the `MQPUT` call, and is passed back to the application in the `JMSMessageID` field. The IBM MQ `messageId` is a 24-byte binary value, whereas the `JMSMessageID` is a string. The `JMSMessageID` is composed of the binary `messageId` value converted to a sequence of 48 hexadecimal characters, prefixed with the characters `ID:`. JMS provides a hint that can be set to disable the production of message identifiers. This hint is ignored, and a unique identifier is assigned in all cases. Any value that is set into the `JMSMessageID` field before a `send()` is overwritten.

If you do require the ability to specify the `MQMD.MessageID`, you can do this with one of the IBM MQ JMS extensions described in “Reading and writing the message descriptor from an IBM MQ classes for JMS application” on page 962.

### **JMSTimestamp to MQRFH2**

During a `send`, the `JMSTimestamp` field is set according to the JVM's clock. This value is set into the `MQRFH2`. Any value that is set into the `JMSTimestamp` field before a `send()` is overwritten. See also the `JMS_IBM_PutDate` and `JMS_IBM_PutTime` properties.

### **JMSType to MQRFH2**

This string is set into the `MQRFH2` `mcd.Type` field. If it is in URI format, it can also affect `mcd.Set` and `mcd.Fmt` fields.

### **JMSCorrelationID to MQMD.CorrelId, MQRFH2**

The `JMSCorrelationID` can hold one of the following:

#### **A provider specific message ID**

This is a message identifier from a message previously sent or received, and so should be a string of 48 lowercase hexadecimal digits that are prefixed with `ID:`. The prefix is removed, the remaining characters are converted into binary, and then they are set into the `MQMD.CorrelId` field. No `CorrelId` value is encoded in the `MQRFH2`.

#### **A provider-native byte[] value**

The value is copied into the `MQMD.CorrelId` field - padded with nulls, or truncated to 24 bytes if necessary. No `CorrelId` value is encoded in the `MQRFH2`.

#### **An application-specific string**

The value is copied into the `MQRFH2`. The first 24 bytes of the string, in UTF8 format, are written into the `MQMD.CorrelID`.



*Mapping JMS property fields:*

These notes refer to the mapping of JMS property fields in IBM MQ messages.

**JMSXUserID from MQMD UserIdentifier**

JMSXUserID is set on return from send call.

**JMSXAppID from MQMD PutApplName**

JMSXAppID is set on return from send call.

**JMSXGroupID to MQRFH2 (point-to-point)**

For point-to-point messages, the JMSXGroupID is copied into the MQMD GroupID field. If the JMSXGroupID starts with the prefix ID:, it is converted into binary. Otherwise, it is encoded as a UTF8 string. The value is padded or truncated if necessary to a length of 24 bytes. The MQMF\_MSG\_IN\_GROUP flag is set.

**JMSXGroupID to MQRFH2 (publish/subscribe)**

For publish/subscribe messages, the JMSXGroupID is copied into the MQRFH2 as a string.

**JMSXGroupSeq MQMD MsgSeqNumber (point-to-point)**

For point-to-point messages, the JMSXGroupSeq is copied into the MQMD MsgSeqNumber field. The MQMF\_MSG\_IN\_GROUP flag is set.

**JMSXGroupSeq MQMD MsgSeqNumber (publish/subscribe)**

For publish/subscribe messages, the JMSXGroupSeq is copied into the MQRFH2 as an i4.

*Mapping JMS provider-specific fields:*

The following notes refer to the mapping of JMS provider-specific fields into IBM MQ messages.

**JMS\_IBM\_Report\_XXX to MQMD Report**

A JMS application can set the MQMD Report options, using the following JMS\_IBM\_Report\_XXX properties. The single MQMD is mapped to several JMS\_IBM\_Report\_XXX properties. The application must set the value of these properties to the standard IBM MQ MQRO\_ constants (included in com.ibm.mq.MQC). So, for example, to request COD with full Data, the application must set JMS\_IBM\_Report\_COD to the value CMQC.MQRO\_COD\_WITH\_FULL\_DATA.

**JMS\_IBM\_Report\_Exception**

MQRO\_EXCEPTION or  
MQRO\_EXCEPTION\_WITH\_DATA or  
MQRO\_EXCEPTION\_WITH\_FULL\_DATA

**JMS\_IBM\_Report\_Expiration**

MQRO\_EXPIRATION or  
MQRO\_EXPIRATION\_WITH\_DATA or  
MQRO\_EXPIRATION\_WITH\_FULL\_DATA

**JMS\_IBM\_Report\_COA**

MQRO\_COA or  
MQRO\_COA\_WITH\_DATA or  
MQRO\_COA\_WITH\_FULL\_DATA

**JMS\_IBM\_Report\_COD**

MQRO\_COD or  
MQRO\_COD\_WITH\_DATA or  
MQRO\_COD\_WITH\_FULL\_DATA

**JMS\_IBM\_Report\_PAN**

MQRO\_PAN

**JMS\_IBM\_Report\_NAN**

MQRO\_NAN

**JMS\_IBM\_Report\_Pass\_Msg\_ID**  
MQRO\_PASS\_MSG\_ID

**JMS\_IBM\_Report\_Pass\_Correl\_ID**  
MQRO\_PASS\_CORREL\_ID

**JMS\_IBM\_Report\_Discard\_Msg**  
MQRO\_DISCARD\_MSG

**JMS\_IBM\_MsgType to MQMD MsgType**

Value maps directly onto MQMD MsgType. If the application has not set an explicit value of JMS\_IBM\_MsgType, a default value is used. This default value is determined as follows:

- If JMSReplyTo is set to an IBM MQ queue destination, MSGType is set to the value MQMT\_REQUEST
- If JMSReplyTo is not set, or is set to anything other than an IBM MQ queue destination, MsgType is set to the value MQMT\_DATAGRAM

**JMS\_IBM\_Feedback to MQMD Feedback**

Value maps directly onto MQMD Feedback.

**JMS\_IBM\_Format to MQMD Format**

Value maps directly onto MQMD Format.

**JMS\_IBM\_Encoding to MQMD Encoding**

If set, this property overrides the numeric encoding of the Destination Queue or Topic.

**JMS\_IBM\_Character\_Set to MQMD CodedCharacterSetId**

If set, this property overrides the coded character set property of the Destination Queue or Topic.

**JMS\_IBM\_PutDate from MQMD PutDate**

The value of this property is set, during send, directly from the PutDate field in the MQMD. Any value that is set into the JMS\_IBM\_PutDate property before a send is overwritten. This field is a String of eight characters, in the IBM MQ Date format of YYYYMMDD. This property can be used with the JMS\_IBM\_PutTime property to determine the time the message was put according to the queue manager.

**JMS\_IBM\_PutTime from MQMD PutTime**

The value of this property is set, during send, directly from the PutTime field in the MQMD. Any value that is set into the JMS\_IBM\_PutTime property before a send is overwritten. This field is a String of eight characters, in the IBM MQ Time format of HHMMSSSTH. This property can be used with the JMS\_IBM\_PutDate property to determine the time the message was put according to the queue manager.

**JMS\_IBM\_Last\_Msg\_In\_Group to MQMD MsgFlags**

For point-to-point messaging, this Boolean value maps to the MQMF\_LAST\_MSG\_IN\_GROUP flag in the MQMD MsgFlags field. It is normally used with the JMSXGroupID and JMSXGroupSeq properties to indicate to a legacy IBM MQ application that this message is the last in a group. This property is ignored for publish/subscribe messaging.

*Mapping IBM MQ fields onto JMS fields (incoming messages):*

These tables show how JMS header and property fields are mapped into MQMD and MQRFH2 fields at get() or receive() time.

Table 126 shows how JMS header fields are mapped onto MQMD/MQRFH2 fields at get() or receive() time. Table 127 shows how JMS property fields are mapped onto MQMD/MQRFH2 fields at get() or receive() time. Table 128 on page 892 shows how JMS provider-specific properties are mapped.

*Table 126. Incoming message JMS header field mapping*

JMS header field name	MQMD field retrieved from	MQRFH2 field retrieved from
JMSDestination		jms.Dst or mqps.Top <sup>1</sup>
JMSDeliveryMode	Persistence <sup>2</sup>	jms.Dlv <sup>2</sup>
JMSExpiration		jms.Exp
JMSPriority	Priority	
JMSMessageID	MsgID	
JMSTimestamp	PutDate <sup>2</sup> PutTime <sup>2</sup>	jms.Tms <sup>2</sup>
JMSCorrelationID	CorrelId <sup>2</sup>	jms.Cid <sup>2</sup>
JMSReplyTo	ReplyToQ <sup>2</sup> ReplyToQMgr <sup>2</sup>	jms.Rto <sup>2</sup>
JMSType		mcd.Type, mcd.Set, mcd.Fmt
JMSRedelivered	BackoutCount	

**Note:**

1. If both jms.Dst and mqps.Top are set, the value in jms.Dst is used.
2. For properties that can have values retrieved from the MQRFH2 or the MQMD, if both are available, the setting in the MQRFH2 is used.
3. JMS\_IBM\_Character\_Set property value is a String value that contains the Java character set equivalent for the numeric CodedCharacterSetId value.

*Table 127. Incoming message property mapping*

JMS property name	MQMD field retrieved from	MQRFH2 field retrieved from
JMSXUserID	UserIdentifier	
JMSXAppID	PutApplName	
JMSXDeliveryCount	BackoutCount	
JMSXGroupID	GroupId <sup>1</sup>	jms.Gid <sup>1</sup>
JMSXGroupSeq	MsgSeqNumber <sup>1</sup>	jms.Seq <sup>1</sup>

**Note:**

1. For properties that can have values retrieved from the MQRFH2 or the MQMD, if both are available, the setting in the MQRFH2 is used. The properties are set from the MQMD values only if the MQMF\_MSG\_IN\_GROUP or MQMF\_LAST\_MSG\_IN\_GROUP message flags are set.

Table 128. Incoming message provider-specific JMS property mapping

JMS property name	MQMD field retrieved from	MQRFH2 field retrieved from
JMS_IBM_Report_Exception	Report	
JMS_IBM_Report_Expiration	Report	
JMS_IBM_Report_COA	Report	
JMS_IBM_Report_COD	Report	
JMS_IBM_Report_PAN	Report	
JMS_IBM_Report_NAN	Report	
JMS_IBM_Report_Pass_Msg_ID	Report	
JMS_IBM_Report_Pass_Correl_ID	Report	
JMS_IBM_Report_Discard_Msg	Report	
JMS_IBM_MsgType	MsgType	
JMS_IBM_Feedback	Feedback	
JMS_IBM_Format	Format	
JMS_IBM_PutApplType	PutApplType	
JMS_IBM_Encoding <sup>1</sup>	Encoding	
JMS_IBM_Character_Set <sup>1</sup>	CodedCharacterSetId	
JMS_IBM_PutDate	PutDate	
JMS_IBM_PutTime	PutTime	
JMS_IBM_Last_Msg_In_Group	MsgFlags	

1. Only set if the incoming message is a Bytes Message.

*Exchanging messages between a JMS application and a traditional IBM MQ application:*

This topic describes what happens when a JMS application exchanges messages with a traditional IBM MQ application that cannot process the MQRFH2 header.

Figure 109 on page 893 shows the mapping.

The administrator indicates that the JMS application is communicating with a traditional IBM MQ application by setting the TARGCLIENT property of the destination to MQ. This indicates that no MQRFH2 header is to be produced. If this is not done, the receiving application must be able to handle the MQRFH2 header.

The mapping from JMS to MQMD targeted at a traditional IBM MQ application is the same as mapping from JMS to MQMD targeted at a JMS application. If IBM MQ classes for JMS receives an IBM MQ message with the MQMD *Format* field set to anything other than MQFMT\_RFH2, data is being received from a non-JMS application. If the format is MQFMT\_STRING, the message is received as a JMS text message. Otherwise, it is received as a JMS bytes message. Because there is no MQRFH2, only those JMS properties that are transmitted in the MQMD can be restored.

If IBM MQ classes for JMS receives a message that does not have an MQRFH2 header, the TARGCLIENT property of the Queue or Topic object derived from the JMSReplyTo header field of the message is set to MQ by default. This means that a reply message sent to the queue or topic also does not have an MQRFH2 header. You can switch off this behavior of including an MQRFH2 header in a reply message only if the original message has an MQRFH2 header, by setting the TARGCLIENTMATCHING property of the connection factory to NO.

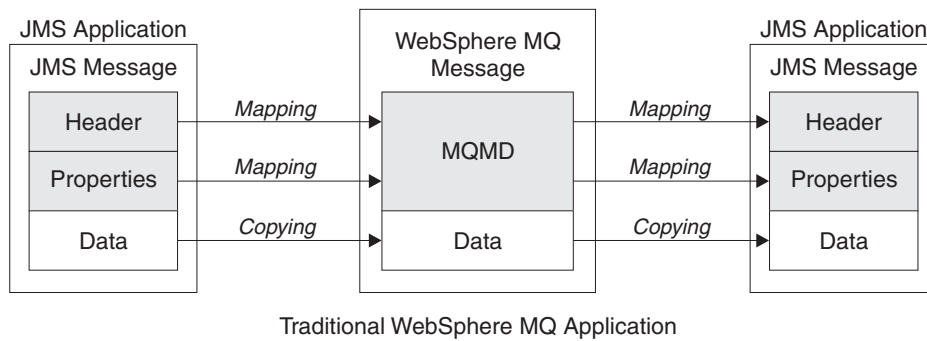


Figure 109. How JMS messages are transformed to IBM MQ messages with no MQRFH2 header

The JMS message body:

This topic contains information about the encoding of the message body itself. The encoding depends on the type of JMS message.

### ObjectMessage

An ObjectMessage is an object serialized by the Java Runtime in the normal way.

### TextMessage

A TextMessage is an encoded string. For an outgoing message, the string is encoded in the character set given by the destination object. This defaults to UTF8 encoding (the UTF8 encoding starts with the first character of the message; there is no length field at the start). It is, however, possible to specify any other character set supported by IBM MQ classes for JMS. Such character sets are used mainly when you send a message to a non-JMS application.

If the character set is a double-byte set (including UTF16), the destination object's integer encoding specification determines the order of the bytes.

An incoming message is interpreted using the character set and encoding that are specified in the message itself. These specifications are in the last IBM MQ header (or MQMD if there are no headers). For JMS messages, the last header is usually the MQRFH2.

### BytesMessage

A BytesMessage is, by default, a sequence of bytes as defined by the JMS 1.0.2 specification and associated Java documentation.

For an outgoing message that was assembled by the application itself, the destination object's encoding property can be used to override the encodings of integer and floating point fields contained in the message. For example, you can request that floating point values are stored in S/390 rather than IEEE format).

An incoming message is interpreted using the numeric encoding specified in the message itself. This specification is in the last IBM MQ header (or MQMD if there are no headers). For JMS messages, the last header is usually the MQRFH2.

If a BytesMessage is received, and is re-sent without modification, its body is transmitted byte for byte, as it was received. The destination object's encoding property has no effect on the body. The only string-like entity that can be sent explicitly in a BytesMessage is a UTF8 string. This is encoded in Java UTF8 format, and starts with a 2-byte length field. The destination object's character set property has no effect on the encoding of an outgoing BytesMessage. The character set value in an incoming IBM MQ message has no effect on the interpretation of that message as a JMS BytesMessage.

Non-Java applications are unlikely to recognize the Java UTF8 encoding. Therefore, for a JMS application to send a `BytesMessage` that contains text data, the application itself must convert its strings to byte arrays, and write these byte arrays into the `BytesMessage`.

### MapMessage

A `MapMessage` is a string containing XML name/type/value triplets encoded as:

```
<map>
 <elt name="elementname1" dt="datatype1">value1</elt>
 <elt name="elementname2" dt="datatype2">value2</elt>
 ...
</map>
```

where `datatype` is one of the data types listed in Table 119 on page 883. The default data type is string, and so the attribute `dt="string"` is omitted for string elements.

The character set used to encode or interpret the XML string that forms the body of a map message is determined according to the rules that apply to a text message.

Versions of IBM MQ classes for JMS earlier than Version 5.3 encoded the body of a map message in the following format:

```
<map>
 <elementname1 dt="datatype1">value1</elementname1>
 <elementname2 dt="datatype2">value2</elementname2>
 ...
</map>
```

Version 5.3 and later versions of IBM MQ classes for JMS can interpret either format, but versions of IBM MQ classes for JMS earlier than Version 5.3 cannot interpret the current format.

If an application needs to send map messages to another application that is using a version of IBM MQ classes for JMS earlier than Version 5.3, the sending application must call the connection factory method `setMapNameStyle(WMQConstants.WMQ_MAP_NAME_STYLE_COMPATIBLE)` to specify that the map messages are sent in the previous format. By default, all map messages are sent in the current format.

### StreamMessage

A `StreamMessage` is like a map message, but without element names:

```
<stream>
 <elt dt="datatype1">value1</elt>
 <elt dt="datatype2">value2</elt>
 ...
</stream>
```

where `datatype` is one of the data types listed in Table 119 on page 883. The default data type is string, and so the attribute `dt="string"` is omitted for string elements.

The character set used to encode or interpret the XML string that makes up the `StreamMessage` body is determined following the rules that apply to a `TextMessage`.

The `MQRFH2.format` field is set as follows:

#### **MQFMT\_NONE**

for `ObjectMessage`, `BytesMessage`, or messages with no body.

#### **MQFMT\_STRING**

for `TextMessage`, `StreamMessage`, or `MapMessage`.

### *JMS message conversion:*

Message data conversion in JMS is performed when sending and receiving messages. IBM MQ performs most data conversion automatically. It converts text and numeric data when transferring a message between JMS applications. Text is converted when exchanging a `JMSTextMessage` between a JMS application and an IBM MQ application.

If you are planning to do more complex message exchanges, the following topics are of interest you. Complex message exchanges include:

- Transferring non-text messages between an IBM MQ application and a JMS application.
- Exchanging text data in byte format.
- Converting the text in your application.

### **JMS message data**

Data conversion is necessary to exchange text and numeric data between applications, even between two JMS applications. The internal representation of text and numbers must be encoded so they can be transferred in a message. Encoding forces a decision about how numbers and text are represented. IBM MQ manages the encoding of text and numbers in JMS messages, except for `JMSObjectMessage`, see “`JMSObjectMessage`” on page 902. It uses three message attributes. The three attributes are `CodedCharacterSetId`, `Encoding`, and `Format`.

These three message attributes are normally stored in the JMS header, `MQRFH2`, fields of a JMS message. If the message type is an MQ, rather than JMS type of message, the attributes are stored in the message descriptor, `MQMD`. The attributes are used to convert the JMS message data. JMS message data is transferred in the message data part of an IBM MQ message.

### **JMS message properties**

JMS message properties, such as `JMS_IBM_CHARACTER_SET`, are exchanged in the `MQRFH2` header part of a JMS message, unless the message has been sent without an `MQRFH2`. Only `JMSTextMessage` and `JMSBytesMessage` can be sent without an `MQRFH2`. If a JMS property is stored as an IBM MQ message property in the message descriptor, `MQMD`, it is converted as part of the `MQMD` conversion. If a JMS property is stored in the `MQRFH2`, it is stored in the character set specified by `MQRFH2.NameValueCCSID`. When a message is sent or received, message properties are converted to and from their internal representation in the JVM. The conversion is to and from the character set of the message descriptor or `MQRFH2.NameValueCCSID`. Numeric data is converted to text.

### **JMS message conversion**

The following topics contain examples and tasks that are useful if you plan to exchange more complex messages that require conversion.

## *JMS message conversion approaches:*

A number of data conversion approaches are open to JMS application designers. These approaches are not exclusive; some applications are likely to use a combination of these approaches. If your application is exchanging only text or is exchanging messages only with other JMS applications, you do not normally consider data conversion. Data conversion is performed automatically for you, by IBM MQ.

You can ask a number of questions about how to approach message conversion:

### **Is it necessary to think about message conversion at all?**

In some cases, such as JMS to JMS message transfers, and exchanging text messages with IBM MQ programs, IBM MQ performs the necessary conversions for you, automatically. You might want to control data conversion for performance reasons, or you might be exchanging complex messages that have a predefined format. In cases such as these you must understand message conversion, and read the following topics.

### **What kinds of conversion are there?**

There are four main types of conversion, which are explained in the following sections:

1. "JMS client data conversion"
2. "Application data conversion" on page 897
3. "Queue manager data conversion" on page 897
4. "Message channel data conversion" on page 898

### **Where should conversion be performed?**

The section, "Choosing an approach to message conversion: "receiver makes good"" on page 899, describes the usual approach of "receiver makes good". "Receiver makes good" also applies to JMS data conversion.

## **JMS client data conversion**

JMS client<sup>4</sup> data conversion is the conversion of Java primitives and objects into bytes in a JMS message as it is sent to a destination, and conversion back again, when it is received. JMS client data conversion uses the methods of the JMSMessage classes. The methods are listed by JMSMessage class type in Table 129 on page 900.

Conversion to and from the internal JVM representation of numbers and text is performed for read, get, set, and write methods. The conversion is performed when the message is sent, and when any of the read or get methods is called on a message that has been received.

The code page and numeric encoding used to write or set the contents of a message are defined as attributes of the destination. The destination code page and numeric encoding can be changed administratively. An application can also override the destination code page and encoding by setting the message properties that control writing or setting message content.

If you want to convert number encoding when a JMSBytesMessage message is sent to a destination that is not defined as Native encoding, you must set the message property JMS\_IBM\_ENCODING before sending the message. If you are following the "receiver makes good" pattern, or if you are exchanging messages between JMS applications, the application does not need to set JMS\_IBM\_ENCODING. In most cases you can leave the Encoding property as Native.

For JMSStreamMessage, JMSMapMessage, and JMSTextMessage messages, the character set identifier properties of the destination are used. Encoding is ignored on sending as numbers are written out in text

---

4. "JMS Client" refers to the IBM MQ classes for JMS that implement the JMS interface, which runs either in client or bindings mode.



format. The JMS client application program does not have to set `JMS_IBM_CHARACTER_SET` before sending the message if the destination character set property to apply.

To get the data in a message an application calls the JMS message read or get methods. The methods refer to the code page and encoding defined in the previous message header to create the Java primitives and objects correctly.

JMS client data conversion meets the needs of most JMS applications that are exchanging messages between one JMS client and another. You do not code any explicit data conversion. You do not use the `java.nio.charset.Charset` class, which is typically used when writing text to a file. The `writeString` and `setString` methods do the conversion for you.

For more details on JMS client data conversion, see “JMS client message conversion and encoding” on page 908.

### Application data conversion

A JMS client application can perform explicit character data conversion by using the `java.nio.charset.Charset` class; see the examples in Figure 112 on page 901 and Figure 113 on page 901. String data is converted into bytes, using the `getBytes` method, and sent as bytes. The bytes are converted back into text by using a `String` constructor that takes a byte array and a `Charset`. Character data is converted using the `encode` and `decode` `Charset` methods. Typically the message is sent or received as `JMSBytesMessage`, because the message part of a `JMSBytesMessage` does not contain anything other than the data written by the application<sup>5</sup>. You can also send and receive bytes using `JMSStreamMessage`, `JMSMapMessage`, or `JMSObjectMessage`.

There are no Java methods to encode and decode bytes that contain numeric data represented in different encoding formats. Numeric data is encoded and decoded automatically using the numeric `JMSMessage` read and write methods. The read and write methods use the value of the `JMS_IBM_ENCODING` attribute of the message data.

A typical use for application data conversion is if a JMS client sends or receives a formatted message from a non-JMS application. A formatted message contains text, numeric, and bytes data organized by the length of the data fields. Unless the non-JMS application has specified the message format as “MQSTR ”, the message is constructed as a `JMSBytesMessage`. To receive formatted message data in a `JMSBytesMessage` you must call a sequence of methods. The methods must be called in the same order the fields were written into the message. If the fields are numeric, you must know the encoding and length of the numeric data. If any of the fields contain byte or text data, you must know the length of any byte data in the message. There are two ways to convert a formatted message into a Java object that is easy to use.

1. Construct a Java class corresponding to the record, to encapsulate reading and writing the message. Access to the data in the record is with `get` and `set` methods of the class.
2. Construct a Java class corresponding to the record by extending the `com.ibm.mq.headers` class. Access to the data in the class is with type-specific accessors of the form, `getStringValue(fieldName)`;

See “Exchanging a formatted record with a non-JMS application” on page 915.

### Queue manager data conversion

In IBM MQ V7.0, code page conversion can be performed by the queue manager when a JMS client program gets a message. The conversion is the same as the conversion performed for a C program. A C program sets `MQGMO_CONVERT` as an `MQGET GetMsgOpts` parameter option; see Figure 111 on page 901. A queue manager performs conversion for a JMS client program that is receiving a message, if the

---

5. One exception: Data written using `writeUTF` starts with a 2 byte length field

WMQ\_RECEIVE\_CONVERSION destination property is set to WMQ\_RECEIVE\_CONVERSION\_QMGR, The JMS client program can also set the destination property; see Figure 110.

Before V7.0, conversions were always performed by the JMS client. JMS client data conversion is restricted to converting sequences of numbers and text of type and length known to the JMS client. It cannot convert data structures; see “Exchanging a formatted record with a non-JMS application” on page 915. In V7.0, until fix pack 7.0.1.5, if the conversion can be performed by the queue manager, it is always performed by the queue manager. From 7.0.1.5 onwards, the default conversion behavior reverts to the same as V6.0, and all conversions are performed by the JMS client. From 7.0.1.5, or 7.0.1.4 with APAR IC72897, you can set a new destination option, WMQ\_RECEIVE\_CONVERSION, to control where conversion is performed, and WMQ\_RECEIVE\_CCSD, to set the target code page; see Figure 110.

---

```
((MQDestination)destination).setIntProperty(
 WMQConstants.WMQ_RECEIVE_CONVERSION,
 WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

Or,

```
((MQDestination)destination).setReceiveConversion
 (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

---

*Figure 110. Enable queue manager data conversion*

The main benefit of queue manager conversion comes when exchanging messages with non-JMS applications. If the Format field in the message is defined, and the target character set, or encoding, is different to the message, the queue manager performs data conversion for the target application, if the application requests it. The queue manager converts message data formatted according to one of the predefined IBM MQ message types, such as a CICS bridge header (MQCIH). If the Format field is user-defined, the queue manager looks for a data conversion exit with the name provided in the Format field.

Queue manager data conversion is used to best effect with the “receiver makes good” design pattern. A sending JMS client does not need to perform conversion. A non-JMS receiving program relies on the conversion exit to ensure that the message is delivered in the required code page and encoding. With a sending JMS client, and non-JMS receiver, the example applies to IBM MQ pre- and post-V7.0. With IBM MQ V7.0, the conversion exit can be called for a receiving JMS program as well.

You can create a data conversion exit, using the data conversion exit utility, **crtmqcvx**, to enable the queue manager to convert your own record formatted data. You can build your own record format, use the `com.ibm.mq.headers` to access it as a Java class, and use your own conversion exit to convert it. On z/OS the utility is called **CSQUCVX**, and on IBM i, **CVTMQMDTA**. See “Exchanging a formatted record with a non-JMS application” on page 915.

### Message channel data conversion

IBM MQ Sender, Server, Cluster-receiver, and Cluster-sender channels have a message conversion option, CONVERT. The contents of a message can optionally be converted when a message is sent. The conversion takes place at the sending end of the channel. The cluster-receiver definition is used to auto-define the corresponding cluster-sender channel.

Data conversion by message channels is typically used if it is not possible to use other forms of conversion.

## Choosing an approach to message conversion: “receiver makes good”

The usual approach in IBM MQ application design for code conversion is “receiver makes good”. “Receiver makes good” reduces the number of message conversions. It also avoids the problem of unexpected channel errors if message conversion fails on some intermediary queue manager during message transfer. The “receiver makes good” rule is only broken if there is some reason why the receiver cannot make good. The receiving platform might not have the right character set, for example.

“Receiver makes good” is also good general guidance for JMS client applications. But in specific cases, conversion to the correct character set at source can be more efficient. Conversion from the JVM internal representation must take place when a message containing text or numeric types is sent. Conversion to the character set required by the receiver, if the receiver is not a JMS client, might remove the need for the non-JMS recipient to perform conversion. If the recipient is a JMS client, it is going to convert again, anyway, to decode the message data and create Java primitives and objects.

The difference between JMS client applications, and applications written in a language such as C, is that Java must perform data conversion. A Java application must convert numbers and text from their internal representation to an encoded format used in messages.

By setting destination, or message properties, you can set the character set and encoding used by IBM MQ to encode numbers and text in messages. Normally, you would leave the character set as `1208` and encoding as `Native`.

IBM MQ does not convert byte arrays. To encode strings and character arrays into byte arrays use the `java.nio.charset` package. `Charset` specifies the character set used to convert a string or character array into a byte array. You can also decode a byte array into a string or character array using a `Charset`. It is not good practice to rely on `java.nio.charset.Charset.defaultCodePage` when encoding strings and character arrays. The default `Charset` is typically `windows-1252` on Windows, and `UTF-8` on UNIX. `windows-1252` is a single-byte character set and `UTF-8` is a multi-byte character set.

Generally leave the destination character set and encoding properties at their default values of `UTF-8` and `Native` when exchanging messages with other JMS applications. If you are exchanging messages containing numbers or text with a JMS application, choose one of the `JMSTextMessage`, `JMSStreamMessage`, `JMSMapMessage`, or `JMSObjectMessage` message types that fit your purpose. There are no other conversion tasks to do.

If you are exchanging messages with non-JMS applications that use a record format, it is more complicated. Unless the entire record contains text and can be transferred as a `JMSTextMessage`, you must encode and decode text in the application. Set the destination message type to `MQ`, and use `JMSBytesMessage` to avoid the IBM MQ classes for JMS adding additional header and tagging information to the message data. Use `JMSBytesMessage` methods to write numbers and bytes, and the `Charset` class convert text into byte arrays explicitly. A number of factors might influence your choice of character set:

- Performance: Can you reduce the number of conversions by transforming text into a character set that is used on the largest number of servers?
- Uniformity: Transfer all messages in the same character set.
- Richness: What character sets have all the code points that applications must use?
- Simplicity: Single-byte character sets are simpler to use than variable length and multibyte character sets.

See “Exchanging a formatted record with a non-JMS application” on page 915. for examples of converting messages exchanged with non-JMS applications.

## Examples

## Table of message types and conversion types

Table 129. Message types and conversion types

Message type	Conversion type			
	Text	Numeric	Other	None
JMSObjectMessage				getObject setObject
JMSTextMessage	getText setText			
JMSBytesMessage	readUTF writeUTF	readDouble readFloat readInt readLong readShort readUnsignedShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readUnsignedByte readBytes readChar writeByte writeBytes writeChar
JMSStreamMessage	readString writeString	readDouble readFloat readInt readLong readShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readBytes readChar writeByte writeBytes writeChar
JMSMapMessage	getString setString	getDouble getFloat getInt getLong getShort setDouble setFloat setInt setLong setShort	getBoolean getObject setBoolean setObject	getByte getBytes readChar setByte setBytes setChar

## Calling data conversion from a C program

---

```

gmo.Options = MQGMO_WAIT /* wait for new messages */
 | MQGMO_NO_SYNCPOINT /* no transaction */
 | MQGMO_CONVERT; /* convert if necessary */
while (CompCode != MQCC_FAILED) {
 buflen = sizeof(buffer) - 1; /* buffer size available for GET */
 memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
 memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
 md.Encoding = MQENC_NATIVE;
 md.CodedCharSetId = MQCCSI_Q_MGR;
 MQGET(Hcon, /* connection handle */
 Hobj, /* object handle */
 &md, /* message descriptor */
 &gmo, /* get message options */
 buflen, /* buffer length */
 buffer, /* message buffer */
 &messlen, /* message length */
 &CompCode, /* completion code */
 &Reason); /* reason code */

```

---

Figure 111. Code snippet from *amqsget0.c*

### Sending and receiving text in a `JMSBytesMessage`

The code in Figure 112 sends a string in a `BytesMessage`. For simplicity, the example sends a single string, for which a `JMSTextMessage` is more appropriate. To receive a text string in bytes message containing a mixture of types, you must know the length of the string in bytes, called `TEXT_LENGTH` in Figure 113. Even for a string with a fixed number of characters, the length of the byte representation might be longer.

---

```

BytesMessage bytes = session.createBytesMessage();
String codePage = CCSID.getCodePage(((MQDestination) destination)
 .getIntProperty(WMQConstants.WMQ_CCSID));
bytes.writeBytes("In the destination code page".getBytes(codePage));
producer.send(bytes);

```

---

Figure 112. Sending a String in a `JMSBytesMessage`

---

```

BytesMessage message = (BytesMessage)consumer.receive();
int TEXT_LENGTH = new Long(message.getBodyLength()).intValue();
byte[] textBytes = new byte[TEXT_LENGTH];
message.readBytes(textBytes, TEXT_LENGTH);
String codePage = message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET);
String textString = new String(textBytes, codePage);

```

---

Figure 113. Receiving a String from a `JMSBytesMessage`

*JMS message types and conversion:*

The choice of message type affects your approach to message conversion. The interaction of message conversion and message type is described for the JMS message types, `JMSObjectMessage`, `JMSTextMessage`, `JMSMapMessage`, `JMSStreamMessage`, and `JMSBytesMessage`.

### **JMSObjectMessage**

`JMSObjectMessage` contains one object, and any objects that it references, serialized into a byte stream by the JVM. Text is serialized into UTF-8, and limited to strings or character arrays of no more than 65534 bytes. An advantage of `JMSObjectMessage` is that applications are not involved in any data conversion issues as long as they use only the methods and attributes of the object. `JMSObjectMessage` provides data conversion for complex objects without the application programmer considering how to encode an object in a message. The disadvantage of using `JMSObjectMessage` is it can be exchanged only with other JMS applications. By choosing one of the other JMS message types, it is possible to exchange JMS messages with non-JMS applications.

“Sending and receiving a `JMSObjectMessage`” on page 904 shows a `String` object being exchanged in a message.

A JMS client application can receive a `JMSObjectMessage` only in a message that has a JMS-style body. The destination must specify a JMS style body.

### **JMSTextMessage**

`JMSTextMessage` contains a single text string. When a text message is sent, the text Format is set to "MQSTR", `WMQConstants.MQFMT_STRING`. The `CodedCharacterSetId` of the text is set to the coded character set identifier defined for its destination. The text is encoded into the `CodedCharacterSetId` by IBM MQ. The `CodedCharacterSetId` and `Format` fields are either set in the message descriptor, `MQMD`, or into the JMS fields in an `MQRFH2`. If the message is defined as having an `WMQ_MESSAGE_BODY_MQ` message body style, or the body style is unspecified, but the target destination is `WMQ_TARGET_DEST_MQ`, then the message descriptor fields are set. Otherwise the message has a JMS `RFH2` and the fields are set in the fixed part of the `MQRFH2`.

An application can override the coded character set identifier defined for a destination. It must set the message property `JMS_IBM_CHARACTER_SET` to a coded character set identifier; see the example in “Sending and receiving a `JMSTextmessage`” on page 905.

When the JMS client calls the `consumer.receive` method queue manager conversion is optional. Queue manager conversion is enabled by setting the destination property `WMQ_RECEIVE_CONVERSION` to `WMQ_RECEIVE_CONVERSION_QMGR`. The queue manager converts the text message from the `JMS_IBM_CHARACTER_SET` specified for the message before transferring the message to the JMS client. The character set of the converted message is 1208, UTF-8, unless the destination has a different `WMQ_RECEIVE_CCSID`. The `CodedCharacterSetId` in the message that refers to the `JMSTextMessage` is updated to the target character set ID. The text is decoded from the target character set into Unicode by the `getText` method; see the example in “Sending and receiving a `JMSTextmessage`” on page 905.

A `JMSTextMessage` can be sent in an MQ-style message body, without a JMS `MQRFH2` header. The value of the destination attributes, `WMQ_MESSAGE_BODY` and `WMQ_TARGET_DEST` determine the message body style, unless overridden by the application. The application can override the values set on the destination by calling `destination.setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_MQ)` or `destination.setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ)`.

If you send a `JMSTextMessage` with an MQ style body by sending it to a destination with `WMQ_MESSAGE_BODY` set to `WMQ_MESSAGE_BODY_MQ`, you cannot receive it as a `JMSTextMessage` from the same destination. All messages received from a destination with `WMQ_MESSAGE_BODY` set to `WMQ_MESSAGE_BODY_MQ`

are received as a `JMSBytesMessage`. If you try to receive the message as a `JMSTextMessage` it causes an exception, `ClassCastException: com.ibm.jms.JMSBytesMessage cannot be cast to javax.jms.TextMessage`.

**Note:** Text in a `JMSBytesMessage` is not converted by the JMS client. The client can only receive the text in the message as a byte array. If queue manager conversion is enabled, the text is converted by the queue manager, but the JMS client must still receive it as a byte array in a `JMSBytesMessage`.

It is generally better to use the `WMQ_TARGET_DEST` property to control whether a `JMSTextMessage` is sent with an MQ or JMS body style. You can then receive the message from a destination that has either `WMQ_TARGET_DEST` set to `WMQ_TARGET_DEST_MQ` or `WMQ_TARGET_DEST_JMS`. `WMQ_TARGET_DEST` has no effect on the receiver.

### **JMSMapMessage and JMSStreamMessage**

These two JMS message types are similar. You can read and write primitive types to the messages using methods based on the `DataInputStream` and `DataOutputStream` interfaces; see “Table of message types and conversion types” on page 907. The details are described in “JMS client message conversion and encoding” on page 908. Each primitive is tagged; see “The JMS message body” on page 893.

Numeric data is read and written to the message encoded as XML text. No reference is made to the destination property, `JMS_IBM_ENCODING`. Text data is treated the same way as text in a `JMSTextMessage`. If you were to look at the message contents created by the example in Figure 118 on page 906, all the message data would be in EBCDIC as it was sent with a character set value of 37.

You can send multiple items in a `JMSMapMessage` or `JMSStreamMessage`.

You can retrieve the individual items of data by name from a `JMSMapMessage`, or by position from a `JMSStreamMessage`. Each item is decoded when a `get` or `read` method is called using the `CodedCharacterSetId` value stored in the message. If the method used to retrieve the item returns a different type to the type that was sent, the type is converted. If the type cannot be converted, an exception is thrown. See Class `JMSStreamMessage` for details. The example in “Sending data in a `JMSStreamMessage` and `JMSMapMessage`” on page 906 illustrates type conversion, and getting the `JMSMapMessage` contents out of sequence.

The `MQRFH2.format` field for the `JMSMapMessage` and `JMSStreamMessage` is set to “MQSTR”. If the destination property `WMQ_RECEIVE_CONVERSION` is set to `WMQ_RECEIVE_CONVERSION_QMGR`, the message data is converted by the queue manager before being sent to the JMS client. The `MQRFH2.CodedCharacterSetId` of the message is the `WMQ_RECEIVE_CCSDID` of the destination. The `MQRFH2.Encoding` is `Native`. If `WMQ_RECEIVE_CONVERSION` is `WMQ_RECEIVE_CONVERSION_CLIENT_MSG` the `CodedCharacterSetId` and `Encoding` of the `MQRFH2` is the value set by the sender.

A JMS client application can receive a `JMSMapMessage` or `JMSStreamMessage` only in a message that has a JMS-style body, and from a destination that does not specify an MQ style body.

### **JMSBytesMessage**

A `JMSBytesMessage` can contain multiple primitive types. You can read and write primitive types to the messages using methods based on the `DataInputStream` and `DataOutputStream` interfaces; see “Table of message types and conversion types” on page 907. The details are described in “JMS message types and conversion” on page 902.

The encoding of numeric data in the message is controlled by the value of `JMS_IBM_ENCODING` that is set before writing numeric data to the `JMSBytesMessage`. An application can override the default `Native` encoding defined for `JMSBytesMessage` by setting the message property `JMS_IBM_ENCODING`.

Text data can be read and written in UTF-8 using the `readUTF` and `writeUTF`, or in Unicode using the `readChar` and `writeChar` methods. There are no methods that use `CodedCharacterSetId`. Alternatively, the JMS client can encode and decode text into bytes using the `Charset` class. It transfers the bytes between the JVM and message without the IBM MQ classes for JMS performing any conversion; see “Sending and receiving text in a `JMSBytesMessage`” on page 906.

A `JMSBytesMessage` sent to an MQ application is typically sent in an MQ-style message body, without a JMS MQRFH2 header. If it is sent to a JMS application, the message body style is typically JMS. The value of the destination attributes, `WMQ_MESSAGE_BODY` and `WMQ_TARGET_DEST` determine the message body style, unless overridden by the application. The application can override the values set on the destination by calling `destination.setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_MQ)` or `destination.setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ)`.

If you send a `JMSBytesMessage` with an MQ style body, you can receive the message from a destination that defines either an MQ or a JMS message body style. If you send a `JMSBytesMessage` with a JMS style body, then you must receive the message from a destination that defines a JMS message body style. If you do not, the MQRFH2 is treated as part of the user message data, which might not be what you are expecting.

Whether a message has an MQ or a JMS body style, the way it is received is not affected by setting `WMQ_TARGET_DEST`.

The message might be transformed later, by the queue manager, if a `Format` is supplied for the message data, and queue manager data conversion is enabled. Do not use the `format` field for anything other than specifying the format of the message data, or leave it blank, `MQConstants.MQFMT_NONE`

You can send multiple items in a `JMSBytesMessage`. Each numeric item is converted when the message is sent using the encoding defined for the message.

You can retrieve the individual items of data from `JMSBytesMessage`. Call `read` methods in the same order as the `write` methods were called to create the message. Each numeric item is converted when the message is called using the `Encoding` value stored in the message.

Unlike `JMSMapMessage` and `JMSStreamMessage`, `JMSBytesMessage` contains only data written by the application. No additional data is stored in the message data, such as the XML tags used to define the items in a `JMSMapMessage` and `JMSStreamMessage`. For this reason, use `JMSBytesMessage` to transfer messages formatted for other applications.

Converting between `JMSBytesMessage` and `DataInputStream` and `DataOutputStream` is useful in some applications. Code based on the example, “Reading and writing messages using `DataInputStream` and `DataOutputStream`” on page 906, is necessary to use the `com.ibm.mq.header` package with JMS.

## Examples

### Sending and receiving a `JMSObjectMessage`



---

```
ObjectMessage omo = session.createObjectMessage();
omo.setObject(new String("A string"));
producer.send(omo);
...
ObjectMessage omi = (ObjectMessage)consumer.receive();
System.out.println((String)omi.getObject());
...
A string
```

---

*Figure 114. Sending and receiving a JMSObjectMessage*

### **Sending and receiving a JMSTextmessage**

A text message cannot contain text in different character sets. The example shows text in different character sets, sent in two different messages.

---

```
TextMessage tmo = session.createTextMessage();
tmo.setText("Sent in the character set defined for the destination");
producer.send(tmo);
```

---

*Figure 115. Send text message in the character set defined by the destination*

---

```
TextMessage tmo = session.createTextMessage();
tmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
tmo.setText("Sent in EBCDIC character set 37");
producer.send(tmo);
```

---

*Figure 116. Send text message in ccsid 37*

---

```
TextMessage tmi = (TextMessage)consumer.receive();
System.out.println(tmi.getText());
...
Sent in the character set defined for the destination
```

---

*Figure 117. Receive text message*

## Sending data in a JMSStreamMessage and JMSMapMessage

---

```
StreamMessage smo = session.createStreamMessage();
smo.writeString("256");
smo.writeInt(512);
smo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
producer.send(smo);
...
MapMessage mmo = session.createMapMessage();
mmo.setString("First", "256");
mmo.setInt("Second", 512);
mmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
producer.send(mmo);
...
StreamMessage smi = (StreamMessage)consumer.receive();
System.out.println("Stream: First as float " + smi.readFloat() +
 " Second as String " + smi.readString());
...
Stream: First as float: 256.0, Second as String: 512
...
MapMessage mmi = (MapMessage)consumer.receive();
System.out.println("Map: Second as String " + mmi.getString("Second") +
 " First as double " + mmi.getDouble("First"));
...
Map: Second as String: 512, First as double: 256.0
```

---

Figure 118. Send data in JMSStreamMessage and JMSMapMessage

## Sending and receiving text in a JMSBytesMessage

The code in Figure 119 sends a string in a BytesMessage. For simplicity, the example sends a single string, for which a JMSTextMessage is more appropriate. To receive a text string in bytes message containing a mixture of types, you must know the length of the string in bytes, called *TEXT\_LENGTH* in Figure 120. Even for a string with a fixed number of characters, the length of the byte representation might be longer.

---

```
BytesMessage bytes = session.createBytesMessage();
String codePage = CCSID.getCodepage(((MQDestination) destination)
 .getIntProperty(WMQConstants.WMQ_CCSID));
bytes.writeBytes("In the destination code page".getBytes(codePage));
producer.send(bytes);
```

---

Figure 119. Sending a String in a JMSBytesMessage

```
BytesMessage message = (BytesMessage)consumer.receive();
int TEXT_LENGTH = new Long(message.getBodyLength()).intValue();
byte[] textBytes = new byte[TEXT_LENGTH];
message.readBytes(textBytes, TEXT_LENGTH);
String codePage = message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET);
String textString = new String(textBytes, codePage);
```

---

Figure 120. Receiving a String from a JMSBytesMessage

## Reading and writing messages using DataInputStream and DataOutputStream

The code in Figure 121 on page 907 creates a JMSBytesMessage using a DataOutputStream.

---

```

ByteArrayOutputStream bout = new ByteArrayOutputStream();
DataOutputStream dout = new DataOutputStream(bout);
BytesMessage messageOut = prod.session.createBytesMessage();
// messageOut.setIntProperty(WMQConstants.JMS_IBM_ENCODING,
// ((MQDestination) prod.destination).getIntProperty
// (WMQConstants.WMQ_ENCODING));
int ccsidOut = (((MQDestination) prod.destination).getIntProperty(WMQConstants.WMQ_CCSID));
String codePageOut = CCSID.getCodepage(ccsidOut);
dout.writeInt(ccsidOut);
dout.write(codePageOut.getBytes(codePageOut));
messageOut.writeBytes(bout.toByteArray());
producer.send(messageOut);

```

---

*Figure 121. Send a JMSBytesMessage using a DataOutputStream*

The statement that sets the JMS\_IBM\_ENCODING property is commented out. The statement is valid, if writing directly to a JMSBytesMessage, but has no effect when writing to DataOutputStream. Numbers that are written to the DataOutputStream are encoded in Native encoding. Setting JMS\_IBM\_ENCODING has no effect.

The code in Figure 122 receives a JMSBytesMessage using a DataInputStream.

---

```

static final int ccsidIn_SIZE = (Integer.SIZE)/8;
...
connection.start();
BytesMessage messageIn = (BytesMessage) consumer.receive();
int messageLength = new Long(messageIn.getBodyLength()).intValue();
byte [] bin = new byte[messageLength];
messageIn.readBytes(bin, messageLength);
DataInputStream din = new DataInputStream(new ByteArrayInputStream(bin));
int ccsidIn = din.readInt();
byte [] codePageByte = new byte[messageLength - ccsidIn_SIZE];
din.read(codePageByte, 0, codePageByte.length);
System.out.println("CCSID " + ccsidIn + " code page " + new String(codePageByte,
 messageIn.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET)));

```

---

*Figure 122. Receive a JMSBytesMessage using a DataInputStream*

The code page is printed out using the code page property of the input message data, JMS\_IBM\_CHARACTER\_SET. On input JMS\_IBM\_CHARACTER\_SET is a Java code page and not a numeric coded character set identifier.

## Table of message types and conversion types

*Table 130. Message types and conversion types*

Message type	Conversion type			
	Text	Numeric	Other	None
JMSObjectMessage				getObject setObject
JMSTextMessage	getText setText			

Table 130. Message types and conversion types (continued)

Message type	Conversion type			
	Text	Numeric	Other	None
JMSBytesMessage	readUTF writeUTF	readDouble readFloat readInt readLong readShort readUnsignedShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readUnsignedByte readBytes readChar writeByte writeBytes writeChar
JMSStreamMessage	readString writeString	readDouble readFloat readInt readLong readShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readBytes readChar writeByte writeBytes writeChar
JMSMapMessage	getString setString	getDouble getFloat getInt getLong getShort setDouble setFloat setInt setLong setShort	getBoolean getObject setBoolean setObject	getByte getBytes readChar setByte setBytes setChar

*JMS client message conversion and encoding:*

The methods you use to do JMS client message conversion and encoding are listed, with code examples of each type of conversion.

Conversion and encoding occur when Java primitives or objects are read or written to and from JMS messages. The conversion is called JMS client data conversion to distinguish it from queue manager data conversion and application data conversion. The conversion takes place strictly when data is read from or written to a JMS message. Text is converted to and from the internal 16 bit Unicode representation<sup>6</sup> to the character set used for text in messages. Numeric data is converted to and from Java primitive numeric types to the encoding defined for the message. Whether conversion is performed, and what type of conversion is performed, depends on the JMS message type and the read or write operation.

Table 131 on page 909 categorizes the read and write methods for different JMS message types by the type of conversion performed. The conversions types are described in the text following the table.

6. Some Unicode representation requires more than 16 bits. See a Java SE reference.

Table 131. Message types and conversion types

Message type	Conversion type			
	Text	Numeric	Other	None
JMSObjectMessage				getObject setObject
JMSTextMessage	getText setText			
JMSBytesMessage	readUTF writeUTF	readDouble readFloat readInt readLong readShort readUnsignedShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readUnsignedByte readBytes readChar writeByte writeBytes writeChar
JMSStreamMessage	readString writeString	readDouble readFloat readInt readLong readShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readBytes readChar writeByte writeBytes writeChar
JMSMapMessage	getString setString	getDouble getFloat getInt getLong getShort setDouble setFloat setInt setLong setShort	getBoolean getObject setBoolean setObject	getByte getBytes readChar setByte setBytes setChar

## Text

The default CodedCharacterSetId for a destination is 1208, UTF-8. By default, text is converted from Unicode and sent as a UTF-8 text string. On receive, the text is converted from the coded character set in the message received by the client, into Unicode.

The setText and writeString methods convert text from Unicode into the character set defined for the destination. An application can override the destination character set by setting the message property JMS\_IBM\_CHARACTER\_SET. JMS\_IBM\_CHARACTER\_SET, when sending a message must be a numeric coded character set identifier<sup>7</sup>.

The code snippets in “Sending and receiving a JMSTextmessage” on page 912 send two messages. One is sent in the character set defined for the destination and the other in character set 37, defined by the application.

7. When receiving a message JMS\_IBM\_CHARACTER\_SET is a Java Charset code page name.

The `getText` and `readString` methods convert the text in the message from the character set defined in the message into Unicode. The methods use the code page defined in the message property, `JMS_IBM_CHARACTER_SET`. The code page is mapped from `MQRFH2.CodedCharacterSetId` unless the message is an MQ-type message and has no `MQRFH2`. If the message is a MQ-type message, with no `MQRFH2`, the code page is mapped from `MQMD.CodedCharacterSetId`.

The code snippet in Figure 127 on page 912 receives the message that was sent to the destination. The text in the message is converted from code page IBM037 back into Unicode.

**Note:** A simple way to check that the text is converted to coded character set 37 is to use IBM MQ Explorer. Browse the queue and show the properties of the message before it is retrieved.

Contrast the code snippet in Figure 126 on page 912 with the incorrect code snippet in Figure 123. In the incorrect snippet the text string is converted twice, once by the application, and again by IBM MQ.

---

```
TextMessage tmo = session.createTextMessage();
tmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
tmo.setText(new String("Sent in EBCDIC character set 37".getBytes(CCSID.getCodepage(37))));
producer.send(tmo);
```

---

*Figure 123. Incorrect code page conversion*

The `writeUTF` method converts text from Unicode to 1208, UTF-8. The text string is prefaced with a 2 byte length. The maximum length of the text string is 65534 bytes. The `readUTF` method reads an item in a message written by the `writeUTF` method. It reads exactly the number of bytes written by the `writeUTF` method.

## Numeric

The default numeric encoding for a destination is `Native`. The `Native` encoding constant for Java has the value 273, `x'00000111'`, which is the same for all platforms. On receive, the numbers in the message are correctly transformed into numeric Java primitives. The transformation uses the encoding defined in the message and the type returned by the `read` method.

The `send` method converts numbers that are added to a message by the `set` and `write` into the numeric encoding defined for the destination. The destination encoding can be overridden for a message by an application setting the message property, `JMS_IBM_ENCODING`; for example:

```
message.setIntProperty(WMQConstants.JMS_IBM_ENCODING, WMQConstants.WMQ_ENCODING_INTEGER_REVERSED);
```

The `get` and `read` numeric methods convert numbers in the message from the numeric encoding defined in the message. They convert the numbers to the type that is specified by the `read` or `get` method; see The `ENCODING` property. The methods use the encoding defined in `JMS_IBM_ENCODING`. The encoding is mapped from `MQRFH2.Encoding` unless the message is an MQ-type message and has no `MQRFH2`. If the message is a MQ-type message, with no `MQRFH2`, then the methods use the encoding defined in `MQMD.Encoding`.

The example in Figure 128 on page 912 shows an application encoding a number in the destination format and sending it in a `JMSStreamMessage`. Compare the example in Figure 128 on page 912 to the example in Figure 129 on page 913. The difference is that `JMS_IBM_ENCODING` must be set in a `JMSBytesMessage`.

**Note:** A simple way to check that the number is encoded correctly is to use IBM MQ Explorer. Browse the queue and show the properties of the message before it is consumed.

## Other

The boolean methods encode `true` and `false` as `x'01'` and `x'00'` in a `JMSByteMessage`, `JMSStreamMessage`, and `JMSMapMessage`.

The UTF methods encode and decode Unicode into UTF-8 text strings. The strings are limited to less than 65536 characters, and are preceded by the 2 byte length field.

The Object methods wrap primitive types as objects. Numeric and text types are encoded or converted as if the primitive types had been read or written using the numeric and text methods.

## None

The `readByte`, `readBytes`, `readUnsignedByte`, `writeByte`, and `writeBytes` methods get or put single bytes, or arrays of bytes, between the application and the message without conversion. The `readChar` and `writeChar` methods get and put 2 byte Unicode characters between the application and the message without conversion.

Using the `readBytes` and `writeBytes` methods, the application can perform its own code point conversion, as in “Sending and receiving text in a `JMSBytesMessage`” on page 913.

IBM MQ does not perform any code page conversion in the client as the message is a `JMSBytesMessage`, and because the `readBytes` and `writeBytes` methods are used. Nonetheless, if the bytes represent text, make sure that code page used by the application matches the coded character set of the destination. The message might be converted again by a queue manager conversion exit. Another possibility is that the receiving JMS client program might follow the convention of converting any byte arrays representing text in the message into strings or characters using the `JMS_IBM_CHARACTER_SET` property in the message.

In this example the client uses the destination coded character set for its conversion:

```
bytes.writeBytes("In the destination code page".getBytes(
 CCSID.getCodepage(((MQDestination) destination)
 .getIntProperty(WMQConstants.WMQ_CCSID))));
```

Alternatively, the client might have chosen a code page and then set the corresponding coded character set in the `JMS_IBM_CHARACTER_SET` property of the message. The IBM MQ classes for Java use `JMS_IBM_CHARACTER_SET` to set the `CodedCharacterSetId` field in the JMS properties in the `MQRFH2`, or in the message descriptor, `MQMD`:

```
String codePage = CCSID.getCodepage(37);
message.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, codePage);8
```

If a byte array is written into a `JMSStringMessage` or `JMSMapMessage`, IBM MQ classes for JMS does not perform data conversion, as the bytes are typed as hexadecimal data not as text in the `JMSStringMessage` and `JMSMapMessage`.

If the bytes represent characters in your application, you must take into account what code points to read and write to the message. The code in Figure 124 follows the convention of using the destination coded character set. If you create the string using the default character set for the JVM, the byte contents depend on the platform. A JVM on Windows typically has a default `Charset` of `windows-1252`, and UNIX, `UTF-8`. Interchange between Windows and UNIX does require that you select an explicit code page for exchanging text as bytes.

---

```
StreamMessage smo = producer.session.createStreamMessage();
smo.writeBytes("123".getBytes(CCSID.getCodepage(((MQDestination) destination)
 .getIntProperty(WMQConstants.WMQ_CCSID))));
```

---

Figure 124. Writing bytes representing a string in a `JMSStreamMessage` using the destination character set

## Examples

---

8. `SetStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET, codePage)` currently accepts only numeric character set identifiers.

## Sending and receiving a JMSTextmessage

A text message cannot contain text in different character sets. The example shows text in different character sets, sent in two different messages.

---

```
TextMessage tmo = session.createTextMessage();
tmo.setText("Sent in the character set defined for the destination");
producer.send(tmo);
```

---

*Figure 125. Send text message in the character set defined by the destination*

---

```
TextMessage tmo = session.createTextMessage();
tmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
tmo.setText("Sent in EBCDIC character set 37");
producer.send(tmo);
```

---

*Figure 126. Send text message in ccsid 37*

---

```
TextMessage tmi = (TextMessage)consumer.receive();
System.out.println(tmi.getText());
...
Sent in the character set defined for the destination
```

---

*Figure 127. Receive text message*

---

## Encoding examples

Examples showing a number being sent in the encoding defines for a destination. Notice that you must set the JMS\_IBM\_ENCODING property of a JMSBytesMessage to the value specified for the destination.

---

```
StreamMessage smo = session.createStreamMessage();
smo.writeInt(256);
producer.send(smo);
...
StreamMessage smi = (StreamMessage)consumer.receive();
System.out.println(smi.readInt());
...
256
```

---

*Figure 128. Sending a number using the destination encoding in a JMSStreamMessage*

---



---

```

BytesMessage bmo = session.createBytesMessage();
bmo.writeInt(256);
int encoding = ((MQDestination) (destination)).getIntProperty
 (WMQConstants.WMQ_ENCODING);
bmo.setIntProperty(WMQConstants.JMS_IBM_ENCODING, encoding);
producer.send(bmo);
...
BytesMessage bmi = (BytesMessage)consumer.receive();
System.out.println(bmi.readInt());
...
256

```

---

*Figure 129. Sending a number using the destination encoding in a JMSBytesMessage*

### **Sending and receiving text in a JMSBytesMessage**

The code in Figure 130 sends a string in a BytesMessage. For simplicity, the example sends a single string, for which a JMSTextMessage is more appropriate. To receive a text string in bytes message containing a mixture of types, you must know the length of the string in bytes, called *TEXT\_LENGTH* in Figure 131. Even for a string with a fixed number of characters, the length of the byte representation might be longer.

---

```

BytesMessage bytes = session.createBytesMessage();
String codePage = CCSID.getCodepage(((MQDestination) destination)
 .getIntProperty(WMQConstants.WMQ_CCSID));
bytes.writeBytes("In the destination code page".getBytes(codePage));
producer.send(bytes);

```

---

*Figure 130. Sending a String in a JMSBytesMessage*

---

```

BytesMessage message = (BytesMessage)consumer.receive();
int TEXT_LENGTH = new Long(message.getBodyLength()).intValue();
byte[] textBytes = new byte[TEXT_LENGTH];
message.readBytes(textBytes, TEXT_LENGTH);
String codePage = message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET);
String textString = new String(textBytes, codePage);

```

---

*Figure 131. Receiving a String from a JMSBytesMessage*

### *Queue manager data conversion:*

Queue manager data conversion has always been available to non-JMS applications receiving messages from JMS clients. Since V7.0, JMS clients receiving messages also use queue manager data conversion. From 7.0.1.5, or 7.0.1.4 with APAR IC72897, queue manager data conversion is optional.

The queue manager can convert character and numeric data in message data using the values of `CodedCharacterSetId`, `Encoding`, and `Format` set for the message data. For non-JMS applications the conversion capability has always been available by setting the `GetMessageOption`, `GMO_CONVERT`. The queue manager conversion capability has not been available to a JMS application receiving a message until V7.0.

You can use queue manager conversion, before V7.0, with a JMS client application that sends a message. The JMS client builds a formatted record, sets the `CodedCharacterSetId`, `Encoding`, and `Format` attributes corresponding to the data placed in the message. A non-JMS receiving application reads the message using `GMO_CONVERT`, and causes a user-written data conversion exit to be called. The data conversion exit is a shared library that has the name set in the `Format` field.

Since V7.0, the queue manager is able to convert messages that are sent to JMS clients. From 7.0.0.0 to 7.0.1.4 inclusive, queue manager conversion is always called for JMS clients. From 7.0.1.5, or from 7.0.1.4 with APAR IC72897 applied, queue manager conversion is controlled by setting the destination property, `WMQ_RECEIVE_CONVERSION`, to `WMQ_RECEIVE_CONVERSION_QMGR`, or `WMQ_RECEIVE_CONVERSION_CLIENT_MSG`. `WMQ_RECEIVE_CONVERSION_CLIENT_MSG` is the default setting, matching the behavior of IBM MQ V6.0, which did not support queue manager data conversion for JMS clients. The application can change the destination setting:

---

```
((MQDestination)destination).setIntProperty(
 WMQConstants.WMQ_RECEIVE_CONVERSION,
 WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

Or,

```
((MQDestination)destination).setReceiveConversion
 (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

---

*Figure 132. Enable queue manager data conversion*

Queue manager data conversion for a JMS client takes place when the client calls a `consumer.receive` method. Text data is transformed into UTF-8 (1208) by default. Subsequent read and get methods decode text in the received data from UTF-8, creating Java text primitives in their internal Unicode encoding. UTF-8 is not the only target character set from queue manager data conversion. You can choose a different CCSID by setting the `WMQ_RECEIVE_CCSID` destination property.

An application can also change the destination setting, for example setting it to 437, DOS-US:

---

```
((MQDestination)destination).setIntProperty
(WMQConstants.WMQ_RECEIVE_CCSID, 437);
```

Or,

```
((MQDestination)destination).setReceiveCCSID(437);
```

---

*Figure 133. Set target coded character set for queue manager conversion*

The reason for changing `WMQ_RECEIVE_CCSID` is specialized; the chosen CCSID makes no difference to the text objects created in the JVM. However, some JVMs, on some platforms, might not be able to handle

conversion from the CCSID of text in the message into Unicode. The option gives you a choice of CCSID for any text delivered to the client in the message. Some JMS client platforms have had problems with message text being delivered in UTF-8.

The JMS code is equivalent to the bold text in the C code in Figure 134,

---

```
gmo.Options = MQGMO_WAIT /* wait for new messages */
 | MQGMO_NO_SYNCPOINT /* no transaction */
 | MQGMO_CONVERT; /* convert if necessary */
while (CompCode != MQCC_FAILED) {
 buflen = sizeof(buffer) - 1; /* buffer size available for GET */
 memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
 memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
 md.Encoding = MQENC_NATIVE;
 md.CodedCharSetId = MQCCSI_Q_MGR;
 MQGET(Hcon, /* connection handle */
 Hobj, /* object handle */
 &md, /* message descriptor */
 &gmo, /* get message options */
 buflen, /* buffer length */
 buffer, /* message buffer */
 &messlen, /* message length */
 &CompCode, /* completion code */
 &Reason); /* reason code */
}
```

---

Figure 134. Code snippet from `amqsget0.c`

#### Note:

Queue manager conversion is only performed on the message data that has a known IBM MQ format. MQSTR, or MQCIH are examples of known formats that are predefined. A known format can also be user-defined format, as long as you have supplied a data-conversion exit.

Messages constructed as `JMSTextMessage`, `JMSMapMessage` and `JMSStreamMessage`, have a MQSTR format, and can be converted by the queue manager.

#### Related concepts:

“Invoking the data-conversion exit” on page 365

A data-conversion exit is a user-written exit that receives control during the processing of an MQGET call.

*Exchanging a formatted record with a non-JMS application:*

Follow the steps suggested in this task to design and build a data conversion exit, and a JMS client application that can exchange messages with a non-JMS application using `JMSBytesMessage`. The exchange of a formatted message with a non-JMS application can take place with or without calling a data conversion exit.

#### Before you begin

You might be able to design a simpler solution to exchanging messages with a non-JMS application using a `JMSTextMessage`. Eliminate that possibility before following the steps in this task.

#### About this task

A JMS client is easier to write if it is not involved in the details of formatting JMS messages exchanged with other JMS clients. As long as the message type is `JMSTextMessage`, `JMSMapMessage`,

JMSStreamMessage, or JMSObjectMessage, IBM MQ looks after the details of formatting the message. IBM MQ deals with differences in code pages and numeric encoding on different platforms.

You can use these message types to exchange messages with non-JMS applications. To do so, you must understand how these messages are constructed by IBM MQ classes for JMS. You might be able to modify the non-JMS application to interpret the messages; see “Mapping JMS messages onto IBM MQ messages” on page 879.

An advantage of using one of these message types is the JMS client programming does not depend on the type of application that it is exchanging messages with. A disadvantage is that it might require a modification to another program, and you might not be able to change the other program.

An alternative approach is to write a JMS client application that can deal with existing message formats. Often existing messages are fixed format and contain a mixture of unformatted data, text, and numbers. Use the steps in this task, and the example JMS client in “Writing classes to encapsulate a record layout in a JMSBytesMessage” on page 918, as a starting point for building a JMS client that can exchange formatted records with non-JMS applications.

### Procedure

1. Define the record layout, or use one of the predefined IBM MQ header classes.  
For handling predefined IBM MQ headers, see Handling IBM MQ message headers.  
Figure 135 on page 917 is an example of a user defined, fixed-length record layout, which can be processed by the data conversion utility.
2. Create the data conversion exit.  
Follow the instructions in Writing a data-conversion exit program to write a data conversion exit.  
To try out the example in “Writing classes to encapsulate a record layout in a JMSBytesMessage” on page 918, name the data conversion exit MYRECORD.
3. Write Java classes to encapsulate the record layout, and sending and receiving record. Two approaches you might take are:
  - Write a class that reads and writes the JMSBytesMessage that contains the record; see “Writing classes to encapsulate a record layout in a JMSBytesMessage” on page 918.
  - Write a class extending `com.ibm.mq.header.Header` to define the data structure of the record; see Creating classes for new header types.
4. Decide what coded character set to exchange messages in.  
See Choosing an approach to message conversion: receiver makes good.
5. Configure the destination to exchange MQ-type messages, without a JMS MQRFH2 header.  
Both the sending and receiving destination must be configured to exchange MQ-type messages. You can use the same destination for both sending and receiving.  
The application can override the destination message body property:  

```
((MQDestination)destination).setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_MQ);
```

  
The example in “Writing classes to encapsulate a record layout in a JMSBytesMessage” on page 918 overrides the destination message body property, ensuring an MQ-style message is sent.
6. Test the solution with JMS and non-JMS applications  
Useful tools to test a data conversion exit are:
  - The `amqsgetc0.c` sample program is useful to test receiving a message sent by a JMS client. See the suggested modifications to use the example header, `RECORD.h`, in Figure 136 on page 918. With the modifications, `amqsgetc0.c` receives a message sent by the example JMS client, `TryMyRecord.java`; see “Writing classes to encapsulate a record layout in a JMSBytesMessage” on page 918.
  - The sample IBM MQ browse program, `amqsbcg0.c`, is useful to inspect the contents of the message header, the JMS header, MQRFH2, and the message contents.

- The SupportPac, ih03: WebSphere Message Broker V7-Message display, test & performance utilities, contains a utility, **rfhutil** to display messages and headers.

If you try to receive messages using the modified `amqsgetc0.c` sample program, and get an error with reason code 2080, check whether the message has an MQRFH2. The modifications assume that the message has been sent to a destination that specifies no MQRFH2.

## Examples

---

```
struct RECORD { MQCHAR StrucID[4];
 MQLONG Version;
 MQLONG StructLength;
 MQLONG Encoding;
 MQLONG CodeCharSetId;
 MQCHAR Format[8];
 MQLONG Flags;
 MQCHAR RecordData[32];
};
```

---

*Figure 135. RECORD.h*

- Declare the RECORD.h data structure

```

struct tagRECORD {
 MQCHAR4 StrucId;
 MQLONG Version;
 MQLONG StrucLength;
 MQLONG Encoding;
 MQLONG CCSID;
 MQCHAR8 Format;
 MQLONG Flags;
 MQCHAR32 RecordData;
};
typedef struct tagRECORD RECORD;
typedef RECORD MQPOINTER PRECORD;
RECORD record;
PRECORD pRecord = &(record);

```

- Modify the MQGET call to use RECORD,

1. Before modification:

```

MQGET(Hcon, /* connection handle */
 Hobj, /* object handle */
 &md, /* message descriptor */
 &gmo, /* get message options */
 buflen, /* buffer length */
 buffer, /* message buffer */
 &messlen, /* message length */
 &CompCode, /* completion code */
 &Reason); /* reason code */

```

2. After modification:

```

MQGET(Hcon, /* connection handle */
 Hobj, /* object handle */
 &md, /* message descriptor */
 &gmo, /* get message options */
 sizeof(RECORD), /* buffer length */
 pRecord, /* message buffer */
 &messlen, /* message length */
 &CompCode, /* completion code */
 &Reason); /* reason code */

```

- Change the print statement,

1. From:

```

buffer[messlen] = '\0'; /* add terminator */
printf("message <%s>\n", buffer);

```

2. To:

```

/* buffer[messlen] = '\0'; add terminator */
printf("ccsid <%d>, flags <%d>, message <%32.32s>\n \0",
 md.CodedCharSetId, record.Flags, record.RecordData);

```

Figure 136. Modify *amqsget0.c*

## Related information:

Utility for creating conversion-exit code

*Writing classes to encapsulate a record layout in a JMSBytesMessage:*

The purpose of this task is to explore, by example, how to combine data conversion and a fixed record layout in a JMSBytesMessage. In the task, you create some Java classes to exchange an example record structure in a JMSBytesMessage. You can modify the example to write classes to exchange other record structures.

A `JMSBytesMessage` is the best choice of JMS message type to exchange mixed data type records with non-JMS programs. It has no additional data inserted into the message body by the JMS provider. It is therefore the best choice of message type to use if a JMS client program interoperates with an existing IBM MQ program. The main challenge in using a `JMSBytesMessage` comes with matching the encoding and character set expected by the other program. A solution is to create a class that encapsulates the record. A class that encapsulates reading and writing a `JMSBytesMessage`, for a specific record type, makes it easier to send and receive fixed-format records in a JMS program. By capturing the generic aspects of the interface in an abstract class, much of the solution can be reused for different record formats. Different record formats can be implemented in classes that extend the abstract generic class.

An alternative approach is to extend the `com.ibm.mq.headers.Header` class. The `Header` class has methods, such as `addMQLONG`, to build a record format in a more declarative way. A disadvantage of using the `Header` class is getting and setting attributes uses a more complicated interpretative interface. Both approaches result in much the same amount of application code.

A `JMSBytesMessage` can encapsulate only a single format, in addition to an `MQRFH2`, in one message, unless each record uses the same format, coded character set, and encoding. The format, encoding, and character set of a `JMSBytesMessage` are properties of all of the message following the `MQRFH2`. The example is written on the assumption that a `JMSBytesMessage` contains only one user record.

### Before you begin

1. Your skill level: you must be familiar with Java programming and JMS. No instructions are provided about setting up the Java development environment. It is advantageous to have written a program to exchange a `JMSTextMessage`, `JMSStreamMessage`, or `JMSMapMessage`. You can then see the differences in exchanging a message using a `JMSBytesMessage`.
2. The example requires IBM WebSphere MQ Version 7.0.
3. The example was created using the Java perspective of the Eclipse workbench. It requires JRE 6.0 or higher. You can use the Java perspective in IBM MQ Explorer to develop and run the Java classes. Alternatively, use your own Java development environment.
4. Using IBM MQ Explorer makes setting up the test environment, and debugging, simpler than using command-line utilities.

### About this task

You are guided through creating two classes: `RECORD` and `MyRecord`. Together these two classes encapsulate a fixed-format record. They have methods to get and set attributes. The `get` method reads the record from a `JMSBytesMessage` and the `put` method writes a record to a `JMSBytesMessage`.

The purpose of the task is not to create a production quality class that you can reuse. You might choose to use the examples in the task to get started on your own classes. The purpose of the task is to provide you with guidance notes, primarily about using character sets, formats, and encoding, when using a `JMSBytesMessage`. Each step in creating the classes is explained, and aspects of using `JMSBytesMessage`, which are sometimes overlooked, are described.

The `RECORD` class is abstract and defines some common fields for a user record. The common fields are modeled on the standard IBM MQ header layout of having an eye catcher, a version, and a length field. The encoding, character set, and format fields, found in many IBM MQ headers, are omitted. Another header cannot follow a user-defined format. The `MyRecord` class, which extends the `RECORD` class, does so by literally extending the record with additional user fields. A `JMSBytesMessage`, created by the classes, can be processed by the queue manager data conversion exit.

“Classes used to run example” on page 925 includes a full listing of `RECORD` and `MyRecord`. It also includes listings of the extra “scaffolding” classes to test the `RECORD` and `MyRecord`. The extra classes are:

## TryMyRecord

The main program to test RECORD and MyRecord.

## EndPoint

An abstract class that encapsulates the JMS connection, destination, and session in a single class. Its interface just meets the needs of testing the RECORD and MyRecord classes. It is not an established design pattern for writing JMS applications.

**Note:** The Endpoint class includes this line of code after creating a destination:

```
((MQDestination)destination).setReceiveConversion
 (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

In V7.0, from V7.0.1.5, it is necessary to turn on queue manager conversion. It is disabled by default. In V7.0, up to V7.0.1.4 queue manager conversion is enabled by default, and this line of code causes an error.

## MyProducer and MyConsumer

Classes that extend EndPoint, and create a MessageConsumer and MessageProducer, connected and ready to accept requests.

Together all the classes make up a complete application you can build and experiment with, to understand how to use data conversion in a JMSBytesMessage.

## Procedure

1. Create an abstract class to encapsulate the standard fields in an IBM MQ header, with a default constructor. Later, you extend the class to tailor the header to your requirements.

```
public abstract class RECORD implements Serializable {
 private static final long serialVersionUID = -1616617232750561712L;
 protected final static int UTF8 = 1208;
 protected final static int MQLONG_LENGTH = 4;
 protected final static int RECORD_STRUCT_ID_LENGTH = 4;
 protected final static int RECORD_VERSION_1 = 1;
 protected final String RECORD_STRUCT_ID = "BLNK";
 protected final String RECORD_TYPE = "BLANK ";
 private String structID = RECORD_STRUCT_ID;
 private int version = RECORD_VERSION_1;
 private int structLength = RECORD_STRUCT_ID_LENGTH + MQLONG_LENGTH * 2;
 private int headerEncoding = WMQConstants.WMQ_ENCODING_NATIVE;
 private String headerCharset = "UTF-8";
 private String headerFormat = RECORD_TYPE;

 public RECORD() {
 super();
 }
}
```

### Note:

- a. The attributes, structID to nextFormat, are listed in the order they are laid out in a standard IBM MQ message header.
  - b. The attributes, format, messageEncoding, and messageCharset, describe the header itself, and are not part of the header.
  - c. You must decide whether to store the coded character set identifier or character set of the record. Java uses character sets and IBM MQ messages use coded character set identifiers. The example code uses character sets.
  - d. int is serialized to MQLONG by IBM MQ. MQLONG is 4 bytes.
2. Create the getters and setters for the private attributes.
    - a. Create or generate the getters:



```

 public String getHeaderFormat() { return headerFormat; }
 public int getHeaderEncoding() { return headerEncoding; }
 public String getMessageCharset() { return headerCharset; }
 public int getMessageEncoding() { return headerEncoding; }
 public String getStructID() { return structID; }
 public int getStructLength() { return structLength; }
 public int getVersion() { return version; }

```

- b. Create or generate the setters:

```

 public void setHeaderCharset(String charset) {
 this.headerCharset = charset; }
 public void setHeaderEncoding(int encoding) {
 this.headerEncoding = encoding; }
 public void setHeaderFormat(String headerFormat) {
 this.headerFormat = headerFormat; }
 public void setStructID(String structID) {
 this.structID = structID; }
 public void setStructLength(int structLength) {
 this.structLength = structLength; }
 public void setVersion(int version) {
 this.version = version; }
}

```

3. Create a constructor to create a RECORD instance from a JMSBytesMessage.

```

public RECORD(BytesMessage message) throws JMSEException, IOException,
 MQDataException {
 super();
 setHeaderCharset(message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET));
 setHeaderEncoding(message.getIntProperty(WMQConstants.JMS_IBM_ENCODING));
 byte[] structID = new byte[RECORD_STRUCT_ID_LENGTH];
 message.readBytes(structID, RECORD_STRUCT_ID_LENGTH);
 setStructID(new String(structID, getMessageCharset()));
 setVersion(message.readInt());
 setStructLength(message.readInt());
}

```

**Note:**

- a. The messageCharset and messageEncoding, are captured from the message properties, as they override the values set for the destination. format is not updated. The example does no error checking. If the Record(BytesMessage) constructor is called, it is assumed that the JMSBytesMessage is a RECORD type message. The line “setStructID(new String(structID, getMessageCharset()))” sets the eye catcher.
  - b. The lines of code that complete the method deserialize fields in the message, in order, updating the default values set in the RECORD instance.
4. Create a put method to write the header fields to a JMSBytesMessage.

```

protected BytesMessage put(MyProducer myProducer) throws IOException,
 JMSEException, UnsupportedEncodingException {
 setHeaderEncoding(myProducer.getEncoding());
 setHeaderCharset(myProducer.getCharset());
 myProducer.setMQClient(true);
 BytesMessage bytes = myProducer.session.createBytesMessage();
 bytes.setStringProperty(WMQConstants.JMS_IBM_FORMAT, getHeaderFormat());
 bytes.setIntProperty(WMQConstants.JMS_IBM_ENCODING, getHeaderEncoding());
 bytes.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET,
 myProducer.getCCSID());
 bytes.writeBytes(String.format("%1$-" + RECORD_STRUCT_ID_LENGTH + " ."
 + RECORD_STRUCT_ID_LENGTH + "s", getStructID())
 .getBytes(getMessageCharset()), 0, RECORD_STRUCT_ID_LENGTH);
 bytes.writeInt(getVersion());
 bytes.writeInt(getStructLength());
 return bytes;
}

```

**Note:**

- a. MyProducer encapsulates the JMS Connection, Destination, Session, and MessageProducer in a single class. MyConsumer, used later on, encapsulates the JMS Connection, Destination, Session, and MessageConsumer in a single class.
- b. For a JMSBytesMessage, if the encoding is other than Native, the encoding must be set in the message. The destination encoding is copied to the message encoding attribute, JMS\_IBM\_CHARACTER\_SET, and saved as an attribute of the RECORD class.
  - 1) "setMessageEncoding(myProducer.getEncoding());" calls "((MQDestination) destination).getIntProperty(WMQConstants.WMQ\_ENCODING);" to get the destination encoding.
  - 2) "Bytes.setIntProperty(WMQConstants.JMS\_IBM\_ENCODING, getMessageEncoding());" sets the message encoding.
- c. The character set used to transform text into bytes is obtained from the destination, and saved as an attribute of the RECORD class. It is not set in the message, because it is not used by the IBM MQ classes for JMS when writing a JMSBytesMessage.

"messageCharset = myProducer.getCharset();" calls

```
public String getCharset() throws UnsupportedEncodingException,
 JMSEException {
 return CCSID.getCodepage(getCCSID());
}
```

It gets the Java character set from a coded character set identifier.

"CCSID.getCodepage(ccsid)" is in the package com.ibm.mq.headers. The ccsid is obtained from another method in MyProducer, which queries the destination:

```
public int getCCSID() throws JMSEException {
 return ((MQDestination) destination)
 .getIntProperty(WMQConstants.WMQ_CCSID);
}
```

- d. "myProducer.setMQClient(true);" overrides the destination setting for the client type, forcing it to an IBM MQ MQI client. You might prefer to omit this line of code, as it obscures an administrative configuration error.

"myProducer.setMQClient(true);" calls:

```
((MQDestination) destination).setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ);
if (!getMQDest()) setMQBody();
```

The code has the side-effect of setting the IBM MQ body style to unspecified, if it must override a setting of JMS.

**Note:**

The IBM MQ classes for JMS write the format, encoding, and character set identifier of the message into the message descriptor, MQMD, or into the JMS header, MQRFH2. It depends on whether the message has an IBM MQ style body. Do not set the MQMD fields manually.

A method exists to set the message descriptor properties manually. It uses the JMS\_IBM\_MQMD\_\* properties. You must set the destination property, WMQ\_MQMD\_WRITE\_ENABLED to set the JMS\_IBM\_MQMD\_\* properties:

```
((MQDestination)destination).setMQMDWriteEnabled(true);
```

You must set the destination property, WMQ\_MQMD\_READ\_ENABLED, to read the properties.

Use the JMS\_IBM\_MQMD\_\* only if you take full control over the whole message payload. Unlike the JMS\_IBM\_\* properties, the JMS\_IBM\_MQMD\_\* properties do not control how IBM MQ classes for JMS construct a JMS message. It is possible to create message descriptor properties that conflict with the properties of the JMS message.

- e. The lines of code that completes the method `serialize` the attributes in class as fields in the message.

The string attributes are padded with blanks. The strings are converted to bytes using the character set defined for the record, and truncated to the length of the message fields.

5. Complete the class by adding the imports.

```
package com.ibm.mq.id;
import java.io.IOException;
import java.io.Serializable;
import java.io.UnsupportedEncodingException;
import javax.jms.BytesMessage;
import javax.jms.JMSException;
import com.ibm.mq.constants.MQConstants;
import com.ibm.mq.headers.MQDataException;
import com.ibm.msg.client.wmq.WMQConstants;
```

6. Create a class to extend the `RECORD` class to include additional fields. Include a default constructor.

```
public class MyRecord extends RECORD {
 private static final long serialVersionUID = -370551723162299429L;
 private final static int FLAGS = 1;
 private final static String STRUCT_ID = "MYRD";
 private final static int DATA_LENGTH = 32;
 private final static String FORMAT = "MYRECORD";
 private int flags = FLAGS;
 private String recordData = "ABCDEFGHIJKLMNOPQRSTUVWXYZ012345";

 public MyRecord() {
 super();
 super.setStructID(STRUCT_ID);
 super.setHeaderFormat(FORMAT);
 super.setStructLength(super.getStructLength() + MQLONG_LENGTH
 + DATA_LENGTH);
 }
}
```

**Note:**

- a. The `RECORD` subclass, `MyRecord`, customizes the eye catcher, format, and length of the header.

7. Create or generate the getters and setters.

- a. Create the getters:

```
public int getFlags() { return flags; }
public String getRecordData() { return recordData; } .
```

- b. Create the setters:

```
public void setFlags(int flags) {
 this.flags = flags; }
public void setRecordData(String recordData) {
 this.recordData = recordData; }
}
```

8. Create a constructor to create a `MyRecord` instance from a `JMSBytesMessage`.

```
public MyRecord(BytesMessage message) throws JMSException, IOException,
 MQDataException {
 super(message);
 setFlags(message.readInt());
 byte[] recordData = new byte[DATA_LENGTH];
 message.readBytes(recordData, DATA_LENGTH);
 setRecordData(new String(recordData, super.getMessageCharset()));
}
```

**Note:**

- a. The fields that make up the standard message template are read first by the `RECORD` class.
- b. The `recordData` text is converted to `String` using the character set property of the message.

9. Create a static method to get a message from a consumer and create a new MyRecord instance.

```
public static MyRecord get(MyConsumer myConsumer) throws JMSEException,
 MQDataException, IOException {
 BytesMessage message = (BytesMessage) myConsumer.receive();
 return new MyRecord(message);
}
```

**Note:**

- a. In the example, for brevity, the MyRecord(BytesMessage) constructor is called from the static get method. Typically, you might separate receiving the message from creating a new MyRecord instance.
10. Create a put method to append the customer fields to a JMSBytesMessage containing a message header.

```
public BytesMessage put(MyProducer myProducer) throws JMSEException,
 IOException {
 BytesMessage bytes = super.put(myProducer);
 bytes.writeInt(getFlags());
 bytes.writeBytes(String.format("%1$-" + DATA_LENGTH + "."
 + DATA_LENGTH + "s", getRecordData())
 .getBytes(super.getMessageCharset()), 0, DATA_LENGTH);
 myProducer.send(bytes);
 return bytes;
}
```

**Note:**

- a. The method calls in the code serialize the attributes in the MyRecord class as fields in the message.
- The recordData String attribute is padded with blanks, converted to bytes using the character set defined for the record, and truncated to the length of the RecordData fields.
11. Complete the class by adding the include statements.

```
package com.ibm.mq.id;
import java.io.IOException;
import javax.jms.BytesMessage;
import javax.jms.JMSEException;
import com.ibm.mq.headers.MQDataException;
```

## Results

### Results:

- The results from running the TryMyRecord class:
  - Sending message in coded character set 37, and using a queue manager conversion exit:  
Out flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID 37 MQ true  
Out flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID 37 MQ true  
In flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 273 CCSID UTF-8
  - Sending message in coded character set 37, and not using a queue manager conversion exit:  
Out flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID 37 MQ true  
Out flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID 37 MQ true  
In flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID IBM037
- The results from modifying the TryMyRecord class not to receive the message, and instead receiving it using the modified amqsget0.c sample. The modified sample accepts a formatted record; see Figure 136 on page 918 in “Exchanging a formatted record with a non-JMS application” on page 915.
  - Sending message in coded character set 37, and using a queue manager conversion exit:  
Sample AMQSGET0 start  
ccsid <850>, flags <1>, message <ABCDEFGHIJKLMNOPQRSTUVWXYZ012345>  
no more messages  
Sample AMQSGET0 end

- Sending message in coded character set 37, and not using a queue manager conversion exit:

```
Sample AMQSGET0 start
MQGET ended with reason code 2110
ccsid <37>, flags <1>, message <---+--+ãÃ++△?ÊëËi△?Î△?+ôôõöμρþÚ±=¾¶§>
no more messages
Sample AMQSGET0 end
```

To try out the example and experiment with different code pages and a data conversion exit. Create the Java classes, configure IBM MQ, and run the main program, TryMyRecord ; see Figure 137 on page 926.

1. Configure IBM MQ and JMS to run the example. The instructions are for running the example on Windows.

- a. Create a queue manager

```
crtmqm -sa -u SYSTEM.DEAD.LETTER.QUEUE QM1
strmqm QM1
```

- b. Create a queue

```
echo DEFINE QL('Q1') REPLACE | runmqsc QM1
```

- c. Create a JNDI directory

```
cd c:\
md JNDI-Directory
```

- d. Switch to the JMS bin directory

The JMS Administration program must be run from here. The path is *MQ\_INSTALLATION\_PATH*\java\bin.

- e. Create the following JMS definitions in a file called JMSQM1Q1.txt

```
DEF CF(QM1) PROVIDERVERSION(7) QMANAGER(QM1)
DEF Q(Q1) CCSID(37) ENCODING(RRR) MSGBODY(MQ) QMANAGER(QM1) QUEUE(Q1) TARGCLIENT(MQ) VERSION(7)
END
```

- f. Run the JMSAdmin program to create the JMS resources

```
JMSAdmin < JMSQM1Q1.txt
```

2. You can create, alter, and browse the definitions you have created using IBM MQ Explorer.
3. Run TryMyRecord.

### Classes used to run example

The classes listed in figures Figure 137 on page 926 to Figure 142 on page 930 are also available in a ZIP file; download jm25529\_.zip or jm25529\_.tar.gz.

---

```
package com.ibm.mq.id;
public class TryMyRecord {
 public static void main(String[] args) throws Exception {
 MyProducer producer = new MyProducer();
 MyRecord outrec = new MyRecord();
 System.out.println("Out flags " + outrec.getFlags() + " text "
 + outrec.getRecordData() + " Encoding "
 + producer.getEncoding() + " CCSID " + producer.getCCSID()
 + " MQ " + producer.getMQDest());
 outrec.put(producer);
 System.out.println("Out flags " + outrec.getFlags() + " text "
 + outrec.getRecordData() + " Encoding "
 + producer.getEncoding() + " CCSID " + producer.getCCSID()
 + " MQ " + producer.getMQDest());
 MyRecord inrec = MyRecord.get(new MyConsumer());
 System.out.println("In flags " + inrec.getFlags() + " text "
 + inrec.getRecordData() + " Encoding "
 + inrec.getMessageEncoding() + " CCSID "
 + inrec.getMessageCharset());
 }
}
```

---

*Figure 137. TryMyRecord*

---

```

package com.ibm.mq.id;
import java.io.IOException;
import java.io.Serializable;
import java.io.UnsupportedEncodingException;
import javax.jms.BytesMessage;
import javax.jms.JMSEException;
import com.ibm.mq.constants.MQConstants;
import com.ibm.mq.headers.MQDataException;
import com.ibm.msg.client.wmq.WMQConstants;
public abstract class RECORD implements Serializable {
 private static final long serialVersionUID = -1616617232750561712L;
 protected final static int UTF8 = 1208;
 protected final static int MQLONG_LENGTH = 4;
 protected final static int RECORD_STRUCT_ID_LENGTH = 4;
 protected final static int RECORD_VERSION_1 = 1;
 protected final String RECORD_STRUCT_ID = "BLNK";
 protected final String RECORD_TYPE = "BLANK ";
 private String structID = RECORD_STRUCT_ID;
 private int version = RECORD_VERSION_1;
 private int structLength = RECORD_STRUCT_ID_LENGTH + MQLONG_LENGTH * 2;
 private int headerEncoding = WMQConstants.WMQ_ENCODING_NATIVE;
 private String headerCharset = "UTF-8";
 private String headerFormat = RECORD_TYPE;

 public RECORD() {
 super();
 }

 public RECORD(BytesMessage message) throws JMSEException, IOException,
 MQDataException {
 super();
 setHeaderCharset(message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET));
 setHeaderEncoding(message.getIntProperty(WMQConstants.JMS_IBM_ENCODING));
 byte[] structID = new byte[RECORD_STRUCT_ID_LENGTH];
 message.readBytes(structID, RECORD_STRUCT_ID_LENGTH);
 setStructID(new String(structID, getMessageCharset()));
 setVersion(message.readInt());
 setStructLength(message.readInt());
 }

 public String getHeaderFormat() { return headerFormat; }
 public int getHeaderEncoding() { return headerEncoding; }
 public String getMessageCharset() { return headerCharset; }
 public int getMessageEncoding() { return headerEncoding; }
 public String getStructID() { return structID; }
 public int getStructLength() { return structLength; }
 public int getVersion() { return version; }
 protected BytesMessage put(MyProducer myProducer) throws IOException,
 JMSEException, UnsupportedEncodingException {
 setHeaderEncoding(myProducer.getEncoding());
 setHeaderCharset(myProducer.getCharset());
 myProducer.setMQClient(true);
 BytesMessage bytes = myProducer.session.createBytesMessage();
 bytes.setStringProperty(WMQConstants.JMS_IBM_FORMAT, getHeaderFormat());
 bytes.setIntProperty(WMQConstants.JMS_IBM_ENCODING, getHeaderEncoding());
 bytes.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET,
 myProducer.getCCSID());
 bytes.writeBytes(String.format("%1$-" + RECORD_STRUCT_ID_LENGTH + " "
 + RECORD_STRUCT_ID_LENGTH + "s", getStructID())
 .getBytes(getMessageCharset()), 0, RECORD_STRUCT_ID_LENGTH);
 bytes.writeInt(getVersion());
 bytes.writeInt(getStructLength());
 return bytes;
 }

 public void setHeaderCharset(String charset) {
 this.headerCharset = charset; }
 public void setHeaderEncoding(int encoding) {
 this.headerEncoding = encoding; }
 public void setHeaderFormat(String headerFormat) {
 this.headerFormat = headerFormat; }
 public void setStructID(String structID) {
 this.structID = structID; }
 public void setStructLength(int structLength) {
 this.structLength = structLength; }
 public void setVersion(int version) {
 this.version = version; }
}

```

---

Figure 138. RECORD

---

```

package com.ibm.mq.id;
import java.io.IOException;
import javax.jms.BytesMessage;
import javax.jms.JMSEException;
import com.ibm.mq.headers.MQDataException;
public class MyRecord extends RECORD {
 private static final long serialVersionUID = -370551723162299429L;
 private final static int FLAGS = 1;
 private final static String STRUCT_ID = "MYRD";
 private final static int DATA_LENGTH = 32;
 private final static String FORMAT = "MYRECORD";
 private int flags = FLAGS;
 private String recordData = "ABCDEFGHIJKLMNOPQRSTUVWXYZ012345";

 public MyRecord() {
 super();
 super.setStructID(STRUCT_ID);
 super.setHeaderFormat(FORMAT);
 super.setStructLength(super.getStructLength() + MQLONG_LENGTH
 + DATA_LENGTH);
 }
 public MyRecord(BytesMessage message) throws JMSEException, IOException,
 MQDataException {
 super(message);
 setFlags(message.readInt());
 byte[] recordData = new byte[DATA_LENGTH];
 message.readBytes(recordData, DATA_LENGTH);
 setRecordData(new String(recordData, super.getMessageCharset()));
 }
 public static MyRecord get(MyConsumer myConsumer) throws JMSEException,
 MQDataException, IOException {
 BytesMessage message = (BytesMessage) myConsumer.receive();
 return new MyRecord(message);
 }
 public int getFlags() { return flags; }
 public String getRecordData() { return recordData; } .

 public BytesMessage put(MyProducer myProducer) throws JMSEException,
 IOException {
 BytesMessage bytes = super.put(myProducer);
 bytes.writeInt(getFlags());
 bytes.writeBytes(String.format("%1$-" + DATA_LENGTH + " ."
 + DATA_LENGTH + "s",getRecordData())
 .getBytes(super.getMessageCharset()), 0, DATA_LENGTH);
 myProducer.send(bytes);
 return bytes;
 }
 public void setFlags(int flags) {
 this.flags = flags; }
 public void setRecordData(String recordData) {
 this.recordData = recordData; }
}

```

---

*Figure 139. MyRecord*



```

package com.ibm.mq.id;
import java.io.UnsupportedEncodingException;
import javax.jms.Connection;
import javax.jms.ConnectionFactory;
import javax.jms.Destination;
import javax.jms.JMSEException;
import javax.jms.Session;
import javax.naming.Context;
import javax.naming.InitialContext;
import javax.naming.NamingException;
import com.ibm.mq.headers.CCSID;
import com.ibm.mq.jms.MQDestination;
import com.ibm.msg.client.wmq.WMQConstants;
public abstract class EndPoint {
 public Context ctx;
 public ConnectionFactory cf;
 public Connection connection;
 public Destination destination;
 public Session session;
 protected EndPoint() throws NamingException, JMSEException {
 System.setProperty("java.naming.provider.url", "file:/C:/JNDI-Directory");
 System.setProperty("java.naming.factory.initial",
 "com.sun.jndi.fscontext.RefFSCContextFactory");
 ctx = new InitialContext();
 cf = (ConnectionFactory) ctx.lookup("QM1");
 connection = cf.createConnection();
 destination = (Destination) ctx.lookup("Q1");
 ((MQDestination)destination).setReceiveConversion
 (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
 session = connection.createSession(false, Session.AUTO_ACKNOWLEDGE); }
 protected EndPoint(String cFactory, String dest) throws NamingException,
 JMSEException {
 System.setProperty("java.naming.provider.url", "file:/C:/JNDI-Directory");
 System.setProperty("java.naming.factory.initial",
 "com.sun.jndi.fscontext.RefFSCContextFactory");
 ctx = new InitialContext();
 cf = (ConnectionFactory) ctx.lookup(cFactory);
 connection = cf.createConnection();
 destination = (Destination) ctx.lookup(dest);
 ((MQDestination)destination).setReceiveConversion
 (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
 session = connection.createSession(false, Session.AUTO_ACKNOWLEDGE); }
 public int getCCSID() throws JMSEException {
 return (((MQDestination) destination)
 .getIntProperty(WMQConstants.WMQ_CCSID)); }
 public String getCharset() throws UnsupportedEncodingException,
 JMSEException {
 return CCSID.getCodepage(getCCSID()); }
 public int getEncoding() throws JMSEException {
 return (((MQDestination) destination)
 .getIntProperty(WMQConstants.WMQ_ENCODING)); }
 public boolean getMQDest() throws JMSEException {
 if (((MQDestination) destination).getMessageBodyStyle()
 == WMQConstants.WMQ_MESSAGE_BODY_MQ)
 || (((MQDestination) destination).getMessageBodyStyle()
 == WMQConstants.WMQ_MESSAGE_BODY_UNSPECIFIED)
 && ((MQDestination) destination).getTargetClient()
 == WMQConstants.WMQ_TARGET_DEST_MQ))
 return true;
 else
 return false; }
 public void setCCSID(int ccsid) throws JMSEException {
 ((MQDestination) destination).setIntProperty(WMQConstants.WMQ_CCSID,
 ccsid); }
 public void setEncoding(int encoding) throws JMSEException {
 ((MQDestination) destination).setIntProperty(WMQConstants.WMQ_ENCODING,
 encoding); }
 public void setMQBody() throws JMSEException {
 ((MQDestination) destination)
 .setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_UNSPECIFIED); }
 public void setMQBody(boolean mqbody) throws JMSEException {
 if (mqbody) ((MQDestination) destination)
 .setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_MQ);
 else ((MQDestination) destination)
 .setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_JMS); }
 public void setMQClient(boolean mqclient) throws JMSEException {
 if (mqclient){
 ((MQDestination) destination).setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ);
 if (!getMQDest()) setMQBody();
 }
 else
 ((MQDestination) destination).setTargetClient(WMQConstants.WMQ_TARGET_DEST_JMS); }
}

```

Figure 140. EndPoint

---

```

package com.ibm.mq.id;
import javax.jms.JMSEException;
import javax.jms.Message;
import javax.jms.MessageProducer;
import javax.naming.NamingException;
public class MyProducer extends EndPoint {
 public MessageProducer producer;
 public MyProducer() throws NamingException, JMSEException {
 super();
 producer = session.createProducer(destination); }
 public MyProducer(String cFactory, String dest) throws NamingException,
 JMSEException {
 super(cFactory, dest);
 producer = session.createProducer(destination); }
 public void send(Message message) throws JMSEException {
 producer.send(message); }
}

```

---

*Figure 141. MyProducer*

---

```

package com.ibm.mq.id;
import javax.jms.JMSEException;
import javax.jms.Message;
import javax.jms.MessageConsumer;
import javax.naming.NamingException;
public class MyConsumer extends EndPoint {
 public MessageConsumer consumer;
 public MyConsumer() throws NamingException, JMSEException {
 super();
 consumer = session.createConsumer(destination);
 connection.start(); }
 public MyConsumer(String cFactory, String dest) throws NamingException,
 JMSEException {
 super(cFactory, dest);
 consumer = session.createConsumer(destination);
 connection.start(); }
 public Message receive() throws JMSEException {
 return consumer.receive(); }
}

```

---

*Figure 142. MyConsumer*

---

### **Creating and configuring connection factories and destinations in an IBM MQ classes for JMS application:**

An IBM MQ classes for JMS application can create connection factories and destinations by retrieving them as administered objects from a Java Naming and Directory Interface (JNDI) namespace, by using the IBM JMS extensions, or by using the IBM MQ JMS extensions. An application can also use the IBM JMS extensions or IBM MQ JMS extensions to set the properties of connection factories and destinations.

Connection factories and destinations are starting points in the flow of logic of a JMS application. An application uses a `ConnectionFactory` object to create a connection to a messaging server, and uses a `Queue` or `Topic` object as a target to send messages to or a source from which to receive messages. An application therefore needs to create at least one connection factory and one or more destinations. Having created a connection factory or destination, the application might then need to configure the object by setting one or more of its properties.

In summary, an application can create and configure connection factories and destinations in the following ways:

#### **Using JNDI to retrieve administered objects**

An administrator can use the IBM MQ JMS administration tool or IBM MQ Explorer to create and configure connection factories and destinations as administered objects in a JNDI namespace. An application can then retrieve the administered objects from the JNDI namespace. Having

retrieved an administered object, the application can, if required, set or change one or more of its properties by using either the IBM JMS extensions or the IBM MQ JMS extensions.

### Using the IBM JMS extensions

An application can use the IBM JMS extensions to create connection factories and destinations dynamically at run time. The application first creates a `JmsFactoryFactory` object, and then uses methods of this object to create connection factories and destinations. Having created a connection factory or destination, the application can use methods inherited from the `JmsPropertyContext` interface to set its properties. Alternatively, the application can use a uniform resource identifier (URI) to specify one or more properties of a destination when it creates the destination.

### Using the IBM MQ JMS extensions

An application can also use the IBM MQ JMS extensions to create connection factories and destinations dynamically at run time. The application uses the supplied constructors to create connection factories and destinations. Having created a connection factory or destination, the application can use methods of the object to set its properties. Alternatively, the application can use a URI to specify one or more properties of a destination when it creates the destination.

### Related information:

Configuring JMS resources

*Using JNDI to retrieve administered objects in a JMS application:*

To retrieve administered objects from a Java Naming and Directory Interface (JNDI) namespace, a JMS application must create an initial context and then use the `lookup()` method to retrieve the objects.

Before an application can retrieve administered objects from a JNDI namespace, an administrator must first create the administered objects. The administrator can use the IBM MQ JMS administration tool or IBM MQ Explorer to create and maintain administered objects in a JNDI namespace. For more information, see [Configuring connection factories and destinations in a JNDI namespace](#).

An application server, typically provides its own repository for administered objects and its own tools for creating and maintaining the objects.

To retrieve administered objects from a JNDI namespace, an application must first create an initial context, as shown in the following example:

```
import javax.jms.*;
import javax.naming.*;
import javax.naming.directory.*;
.
.
String url = "ldap://server.company.com/o=company_us,c=us";
String icf = "com.sun.jndi.ldap.LdapCtxFactory";
.
java.util.Hashtable environment = new java.util.Hashtable();
environment.put(Context.PROVIDER_URL, url);
environment.put(Context.INITIAL_CONTEXT_FACTORY, icf);
Context ctx = new InitialDirContext(environment);
```

In this code, the String variables `url` and `icf` have the following meanings:

**url** The uniform resource locator (URL) of the directory service. The URL can have one of the following formats:

- `ldap://hostname/contextName` , for a directory service based on an LDAP server
- `file://directoryPath` , for a directory service based on the local file system

**icf** The class name of the initial context factory, which can be one of the following values:

- `com.sun.jndi.ldap.LdapCtxFactory`, for a directory service based on an LDAP server

- `com.sun.jndi.fscontext.RefFSContextFactory`, for a directory service based on the local file system

Note that some combinations of a JNDI package and a Lightweight Directory Access Protocol (LDAP) service provider can cause LDAP error 84 to occur. To resolve this problem, insert the following line of code before the call to `InitialDirContext()`:

```
environment.put(Context.REFERRAL, "throw");
```

After an initial context is obtained, the application can retrieve administered objects from the JNDI namespace by using the `lookup()` method, as shown in the following example:

```
ConnectionFactory factory;
Queue queue;
Topic topic;
.
.
.
factory = (ConnectionFactory)ctx.lookup("cn=myCF");
queue = (Queue)ctx.lookup("cn=myQ");
topic = (Topic)ctx.lookup("cn=myT");
```

This code retrieves the following objects from an LDAP based namespace:

- A `ConnectionFactory` object bound with the name `myCF`
- A `Queue` object bound with the name `myQ`
- A `Topic` object bound with the name `myT`

For more information about using JNDI, see the JNDI documentation provided by Oracle Corporation.

**Related information:**

Configuring JMS objects using IBM MQ Explorer

Configuring JMS objects using the administration tool

Configuring JMS resources in WebSphere Application Server

*Using the IBM JMS extensions:*

IBM MQ classes for JMS contains a set of extensions to the JMS API called the IBM JMS extensions. An application can use these extensions to create connection factories and destinations dynamically at run time, and to set the properties of IBM MQ classes for JMS objects. The extensions can be used with any messaging provider.

The IBM JMS extensions are a set of interfaces and classes in the following packages:

- `com.ibm.msg.client.jms`
- `com.ibm.msg.client.services`

The packages can be found in `com.ibm.mqjms.jar` which is located in `MQ_INSTALLATION_PATH/java/lib`.

These extensions provide the following function:

- A factory-based mechanism for creating connection factories and destinations dynamically at run time, instead of retrieving them as administered objects from a Java Naming and Directory Interface (JNDI) namespace
- A set of methods for setting the properties of IBM MQ classes for JMS objects
- A set of exception classes with methods for obtaining detailed information about a problem
- A set of methods for controlling tracing
- A set of methods for obtaining version information about IBM MQ classes for JMS

With regard to creating connection factories and destinations dynamically at run time, and setting and getting their properties, the IBM JMS extensions provide an alternative set of interfaces to the IBM MQ JMS extensions. However, whereas the IBM MQ JMS extensions are specific to the IBM MQ messaging provider, the IBM JMS extensions are not specific to IBM MQ and can be used with any messaging provider within the layered architecture described in Layered architecture.

The interface `com.ibm.msg.client.wmq.WMQConstants` contains the definitions of constants, which an application can use when setting the properties of IBM MQ classes for JMS objects using the IBM JMS extensions. The interface contains constants for the IBM MQ messaging provider and JMS constants that are independent of any messaging provider.

The examples of code that follow assume that the following import statements have been run:

```
import com.ibm.msg.client.jms.*;
import com.ibm.msg.client.services.*;
import com.ibm.msg.client.wmq.WMQConstants;
```

### Creating connection factories and destinations

Before an application can create connection factories and destinations using the IBM JMS extensions, it must first create a `JmsFactoryFactory` object. To create a `JmsFactoryFactory` object, the application calls the `getInstance()` method of the `JmsFactoryFactory` class, as shown in the following example:

```
JmsFactoryFactory ff = JmsFactoryFactory.getInstance(JmsConstants.WMQ_PROVIDER);
```

The parameter on the `getInstance()` call is a constant that identifies the IBM MQ messaging provider as the chosen messaging provider. The application can then use the `JmsFactoryFactory` object to create connection factories and destinations.

To create a connection factory, the application calls the `createConnectionFactory()` method of the `JmsFactoryFactory` object, as shown in the following example:

```
JmsConnectionFactory factory = ff.createConnectionFactory();
```

This statement creates a `JmsConnectionFactory` object with the default values for all its properties, which means that the application connects to the default queue manager in bindings mode. If you want an application to connect in client mode, or connect to a queue manager other than the default queue manager, the application must set the appropriate properties of the `JmsConnectionFactory` object before creating the connection. For information about how to do this, see “Setting the properties of IBM MQ classes for JMS objects” on page 934.

The `JmsFactoryFactory` class also contains methods to create connection factories of the following types:

- `JmsQueueConnectionFactory`
- `JmsTopicConnectionFactory`
- `JmsXAConnectionFactory`
- `JmsXAQueueConnectionFactory`
- `JmsXATopicConnectionFactory`

To create a `Queue` object, the application calls the `createQueue()` method of the `JmsFactoryFactory` object, as shown in the following example:

```
JmsQueue q1 = ff.createQueue("Q1");
```

This statement creates an `JmsQueue` object with the default values for all its properties. The object represents an IBM MQ queue called Q1 that belongs to the local queue manager. This queue can be a local queue, an alias queue, or a remote queue definition.

The `createQueue()` method can also accept a queue uniform resource identifier (URI) as a parameter. A queue URI is a string that specifies the name of an IBM MQ queue and, optionally, the name of the queue manager that owns the queue, and one or more properties of the `JmsQueue` object. The following statement contains an example of a queue URI:

```
JmsQueue q2 = ff.createQueue("queue://QM2/Q2?persistence=2&priority=5");
```

The `JmsQueue` object created by this statement represents an IBM MQ queue called Q2 that is owned by queue manager QM2, and all messages sent to this destination are persistent and have a priority of 5. For more information about queue URIs, see “Uniform resource identifiers (URIs)” on page 946. For an alternative way of setting the properties of a `JmsQueue` object, see “Setting the properties of IBM MQ classes for JMS objects.”

To create a `Topic` object, an application can use the `createTopic()` method of the `JmsFactoryFactory` object, as shown in the following example:

```
JmsTopic t1 = ff.createTopic("Sport/Football/Results");
```

This statement creates a `JmsTopic` object with the default values for all its properties. The object represents a topic called `Sport/Football/Results`.

The `createTopic()` method can also accept a topic URI as a parameter. A topic URI is a string that specifies the name of a topic and, optionally, one or more properties of the `JmsTopic` object. The following statements contain an example of a topic URI:

```
String s1 = "topic://Sport/Tennis/Results?persistence=1&priority=0";
JmsTopic t2 = ff.createTopic(s1);
```

The `JmsTopic` object created by these statements represents a topic called `Sport/Tennis/Results`, and all messages sent to this destination are nonpersistent and have a priority of 0. For more information about topic URIs, see “Uniform resource identifiers (URIs)” on page 946. For an alternative way of setting the properties of a `JmsTopic` object, see “Setting the properties of IBM MQ classes for JMS objects.”

After an application has created a connection factory or destination, that object can be used only with the selected messaging provider.

### Setting the properties of IBM MQ classes for JMS objects

To set the properties of IBM MQ classes for JMS objects using the IBM JMS extensions, an application uses the methods of the `com.ibm.msg.client.JmsPropertyContext` interface.

For each Java data type, the `JmsPropertyContext` interface contains a method to set the value of a property with that data type, and a method to get the value of a property with that data type. For example, an application calls the `setIntProperty()` method to set a property with an integer value, and calls the `getIntProperty()` method to get a property with an integer value.

Instances of classes in the `com.ibm.mq.jms` package also inherit the methods of the `JmsPropertyContext` interface. An application can therefore use these methods to set the properties of `MQConnectionFactory`, `MQQueue`, and `MQTopic` objects.

When an application creates an IBM MQ classes for JMS object, any properties with default values are set automatically. When an application sets a property, the new value replaces any previous value the property had. After a property has been set, it cannot be deleted, but its value can be changed.

If an application attempts to set a property to a value that is not valid value for the property, IBM MQ classes for JMS throws a `JMSEException` exception. If an application attempts to get a property that has not been set, the behavior is as described in the JMS specification. IBM MQ classes for JMS throws a `NumberFormatException` exception for primitive data types and returns null for referenced data types.

In addition to the predefined properties of an IBM MQ classes for JMS object, an application can set its own properties. These application defined properties are ignored by IBM MQ classes for JMS.

For more information about the properties of IBM MQ classes for JMS objects, see Properties of IBM MQ classes for JMS objects.

The following code is an example of how to set properties using the IBM JMS extensions. The code sets five properties of a connection factory.

```
factory.setIntProperty(WMQConstants.WMQ_CONNECTION_MODE,
 WMQConstants.WMQ_CM_CLIENT);
factory.setStringProperty(WMQConstants.WMQ_QUEUE_MANAGER, "QM1");
factory.setStringProperty(WMQConstants.WMQ_HOST_NAME, "HOST1");
factory.setIntProperty(WMQConstants.WMQ_PORT, 1415);
factory.setStringProperty(WMQConstants.WMQ_CHANNEL, "QM1.SVR");
factory.setStringProperty(WMQConstants.WMQ_APPLICATIONNAME, "My Application");
```

The effect of setting these properties is that the application connects to queue manager QM1 in client mode, using an MQI channel called QM1.SVR. The queue manager is running on a system with host name HOST1, and the listener for the queue manager is listening in port number 1415. This connection and other queue manager connections associated with sessions under it, have the application name "My Application" associated with them.

**Note:** Queue managers running on z/OS platforms do not support setting application names, and this setting is therefore ignored.

The `JmsPropertyContext` interface also contains the `setObjectProperty()` method, which an application can use to set properties. The second parameter of the method is an object that encapsulates the value of the property. For example, the following code creates an `Integer` object that encapsulates the integer 1415, and then calls `setObjectProperty()` to set the `PORT` property of a connection factory to the value 1415:

```
Integer port = new Integer(1415);
factory.setObjectProperty(WMQConstants.WMQ_PORT, port);
```

This code is therefore equivalent to the following statement:

```
factory.setIntProperty(WMQConstants.WMQ_PORT, 1415);
```

Conversely, the `getObjectProperty()` method returns an object that encapsulates the value of a property.

### Implicit conversion of a property value from one data type to another

When an application uses a method of the `JmsPropertyContext` interface to set or get the property of an IBM MQ classes for JMS object, the value of the property can be implicitly converted from one data type to another.

For example, the following statement sets the `PRIORITY` property of the `JmsQueue` object `q1`:

```
q1.setStringProperty(WMQConstants.WMQ_PRIORITY, "5");
```

The `PRIORITY` property has an integer value, and so the `setStringProperty()` call implicitly converts the string "5" (the source value) to the integer 5 (the target value), which then becomes the value of the `PRIORITY` property.

Conversely, the following statement gets the `PRIORITY` property of the `JmsQueue` object `q1`:

```
String s1 = q1.getStringProperty(WMQConstants.WMQ_PRIORITY);
```

The integer 5 (the source value), which is the value of the `PRIORITY` property, is implicitly converted to the string "5" (the target value) by the `getStringProperty()` call.

The conversions supported by IBM MQ classes for JMS are shown in Table 132.

Table 132. Supported conversions from one data type to another

Source data type	Supported target data types
boolean	String
byte	int, long, short, String
char	String
double	String
float	double, String
int	long, String
long	String
short	int, long, String
String	boolean, byte, double, float, int, long, short

The general rules governing the supported conversions are as follows:

- Numeric values can be converted from one data type to another provided no data is lost during the conversion. For example, a value with data type `int` can be converted into a value with data type `long`, but cannot be converted into a value with data type `short`.
- A value of any data type can be converted into a string.
- A string can be converted to a value of any other data type (except `char`) provided the string is in the correct format for the conversion. If an application attempts to convert a string that is not in the correct format, IBM MQ classes for JMS throws a `NumberFormatException` exception.
- If an application attempts a conversion that is not supported, IBM MQ classes for JMS throws a `MessageFormatException` exception.

The specific rules for converting a value from one data type to another are as follows:

- When converting a boolean value to a string, the value `true` is converted to the string `"true"`, and the value `false` is converted to the string `"false"`.
- When converting a string to a boolean value, the string `"true"` (not case-sensitive) is converted to `true`, and the string `"false"` (not case-sensitive) is converted to `false`. Any other string is converted to `false`.
- When converting a string to a value with data type `byte`, `int`, `long`, or `short`, the string must have the following format:

[ *blanks* ][ *sign* ] *digits*

The meanings of the components of the string are as follows:

*blanks* Optional leading blank characters.

*sign* An optional plus sign (+) or minus sign (-).

*digits* A contiguous sequence of digits (0-9). At least one digit must be present.

After the sequence of digits, the string can contain other characters that are not digits, but the conversion stops as soon as the first of these characters is reached. The string is assumed to represent a decimal integer.

If the string is not in the correct format, IBM MQ classes for JMS throws a `NumberFormatException` exception.

- When converting a string to a value with data type `double` or `float`, the string must have the following format:

[ *blanks* ][ *sign* ] *digits* [ *e\_char* [ *e\_sign* ] *e\_digits* ]

The meanings of the components of the string are as follows:



*blanks* Optional leading blank characters.

*sign* An optional plus sign (+) or minus sign (-).

*digits* A contiguous sequence of digits (0-9). At least one digit must be present.

*e\_char* An exponent character, which is either *E* or *e*.

*e\_sign* An optional plus sign (+) or minus sign (-) for the exponent.

*e\_digits*

A contiguous sequence of digits (0-9) for the exponent. At least one digit must be present if the string contains an exponent character.

After the sequence of digits, or the optional characters representing an exponent, the string can contain other characters that are not digits, but the conversion stops as soon as the first of these characters is reached. The string is assumed to represent a decimal floating point number with an exponent that is a power of 10.

If the string is not in the correct format, IBM MQ classes for JMS throws a `NumberFormatException` exception.

- When converting a numeric value (including a value with data type `byte`) to a string, the value is converted to the string representation of the value as a decimal number, not the string containing the ASCII character for that value. For example, the integer 65 is converted to the string "65", not the string "A".

### Setting more than one property in a single call

The `JmsPropertyContext` interface also contains the `setBatchProperties()` method, which an application can use to set more than one property in a single call. The parameter of the method is a `Map` object that encapsulates a set of property name-value pairs.

For example, the following code uses the `setBatchProperties()` method to set the same five properties of a connection factory as shown in "Setting the properties of IBM MQ classes for JMS objects" on page 934. The code creates an instance of the `HashMap` class, which implements the `Map` interface.

```
HashMap batchProperties = new HashMap();
batchProperties.put(WMQConstants.WMQ_CONNECTION_MODE,
 new Integer(WMQConstants.WMQ_CM_CLIENT));
batchProperties.put(WMQConstants.WMQ_QUEUE_MANAGER, "QM1");
batchProperties.put(WMQConstants.WMQ_WMQ_HOST_NAME, "HOST1");
batchProperties.put(WMQConstants.WMQ_PORT, "1414");
batchProperties.put(WMQConstants.WMQ_CHANNEL, "QM1.SVR");
factory.setBatchProperties(batchProperties);
```

Note that the second parameter of the `Map.put()` method must be an object. Therefore a property value with a primitive data type must be encapsulated within an object or represented by a string, as shown in the example.

The `setBatchProperties()` method validates each property. If the `setBatchProperties()` method cannot set a property because, for example, its value is not valid, none of the specified properties are set.

### Property names and values

If an application uses the methods of the `JmsPropertyContext` interface to set and get the properties of IBM MQ classes for JMS objects, the application can specify the names and values of properties in any of the following ways. Each of the accompanying examples shows how to set the `PRIORITY` property of the `JmsQueue` object `q1` so that a message sent to the queue has the priority specified on the `send()` call.

### Using the property names and values that are defined as constants in the `com.ibm.msg.client.wmq.WMQConstants` interface

The following statement is an example of how to specify the names and values of properties in this way:

```
q1.setIntProperty(WMQConstants.WMQ_PRIORITY, WMQConstants.WMQ_PRI_APP);
```

### Using the property names and values that can be used in queue and topic uniform resource identifiers (URIs)

The following statement is an example of how to specify the names and values of properties in this way:

```
q1.setIntProperty("priority", -2);
```

Only the names and values of properties of destinations can be specified in this way.

### Using the property names and values that are recognized by the IBM MQ JMS administration tool

The following statement is an example of how to specify the names and values of properties in this way:

```
q1.setStringProperty("PRIORITY", "APP");
```

The short form of the property name is also acceptable, as shown in the following statement:

```
q1.setStringProperty("PRI", "APP");
```

When an application gets a property, the value returned depends on the way in which the application specifies the name of the property. For example, if an application specifies the constant `WMQConstants.WMQ_PRIORITY` as the property name, the value returned is the integer -2:

```
int n1 = getIntProperty(WMQConstants.WMQ_PRIORITY);
```

The same value is returned if the application specifies the string "priority" as the property name:

```
int n2 = getIntProperty("priority");
```

However, if the application specifies the string "PRIORITY" or "PRI" as the property name, the value returned is the string "APP":

```
String s1 = getStringProperty("PRI");
```

Internally, IBM MQ classes for JMS stores property names and values as the literal values defined in the `com.ibm.msg.client.wmq.WMQConstants` interface. This is the defined canonical format for property names and values. As a general rule, if an application sets properties using one of the other two ways of specifying property names and values, IBM MQ classes for JMS has to convert the names and values from the specified input format into the canonical format. Similarly, if an application gets properties using one of the other two ways of specifying property names and values, IBM MQ classes for JMS must convert the names from the specified input format into the canonical format, and convert the values from the canonical format into the required output format. Having to perform these conversions might have implications for performance.

Property names and values returned by exceptions, in trace files, or in the IBM MQ classes for JMS log are always in the canonical format.

### Using the Map interface

The `JmsPropertyContext` interface extends the `java.util.Map` interface. An application can therefore use the methods of the `Map` interface to access the properties of an IBM MQ classes for JMS object.

For example, the following code prints out the names and values of all the properties of a connection factory. The code uses only the methods of the `Map` interface to get the names and values of the properties.

```
// Get the names of all the properties
Set propNameNames = factory.keySet();

// Loop round all the property names and get the property values
Iterator iterator = propNameNames.iterator();
while (iterator.hasNext()){
 String propName = (String)iterator.next();
 System.out.println(propName+"="+factory.get(propName));
}
```

Using the methods of the Map interface does not bypass any property validations or conversions.

*Using the IBM MQ JMS extensions:*

IBM MQ classes for JMS contains a set of extensions to the JMS API called the IBM MQ JMS extensions. An application can use these extensions to create connection factories and destinations dynamically at run time, and to set the properties of connection factories and destinations.

IBM MQ classes for JMS contains a set of classes in the packages `com.ibm.jms` and `com.ibm.mq.jms`. These classes implement the JMS interfaces and contain the IBM MQ JMS extensions. The examples of code that follow assume that these packages have been imported by the following statements:

```
import com.ibm.jms.*;
import com.ibm.mq.jms.*;
import com.ibm.msg.client.wmq.WMQConstants;
```

An application can use the IBM MQ JMS extensions to perform the following functions:

- Create connection factories and destinations dynamically at run time, instead of retrieving them as administered objects from a Java Naming and Directory Interface (JNDI) namespace
- Set the properties of connection factories and destinations

### Creating connection factories

To create a connection factory, an application can use the `MQConnectionFactory` constructor, as shown in the following example:

```
MQConnectionFactory factory = new MQConnectionFactory();
```

This statement creates an `MQConnectionFactory` object with the default values for all its properties, which means that the application connects to the default queue manager in bindings mode. If you want an application to connect in client mode, or connect to a queue manager other than the default queue manager, the application must set the appropriate properties of the `MQConnectionFactory` object before creating the connection. For information about how to do this, see “Setting the properties of connection factories.”

An application can create connection factories of the following types in a similar way:

- `MQQueueConnectionFactory`
- `MQTopicConnectionFactory`
- `MQXAConnectionFactory`
- `MQXAQueueConnectionFactory`
- `MQXATopicConnectionFactory`

### Setting the properties of connection factories

An application can set the properties of a connection factory by calling the appropriate methods of the connection factory. The connection factory can either be an administered object or an object created dynamically at run time.

Consider the following code, for example:

```
MQConnectionFactory factory = new MQConnectionFactory();
.
factory.setTransportType(WMQConstants.WMQ_CM_CLIENT);
factory.setQueueManager("QM1");
factory.setHostName("HOST1");
factory.setPort(1415);
factory.setChannel("QM1.SVR");
```

This code creates an `MQConnectionFactory` object and then sets five properties of the object. The effect of setting these properties is that the application connects to queue manager QM1 in client mode using an MQI channel called QM1.SVR. The queue manager is running on a system with host name HOST1, and the listener for the queue manager is listening in port number 1415.

For a real-time connection to a broker, an application can use the following code:

```
MQConnectionFactory factory = new MQConnectionFactory();
.
factory.setTransportType(WMQConstants.WMQ_CM_DIRECT);
factory.setHostName("HOST2");
factory.setPort(1507);
```

This code assumes that the broker is running on a system with host name HOST2 and listening on port number 1507.

An application that uses a real-time connection to a broker can use only the publish/subscribe style of messaging. It cannot use the point-to-point style of messaging.

Only certain combinations of properties of a connection factory are valid. For information about which combinations are valid, see *Dependencies between properties of IBM MQ classes for JMS objects*.

For more information about the properties of a connection factory, and the methods used to set its properties, see *Properties of IBM MQ classes for JMS objects*.

## Creating destinations

To create a `Queue` object, an application can use the `MQQueue` constructor, as shown in the following example:

```
MQQueue q1 = new MQQueue("Q1");
```

This statement creates an `MQQueue` object with the default values for all its properties. The object represents an IBM MQ queue called Q1 that belongs to the local queue manager. This queue can be a local queue, an alias queue, or a remote queue definition.

An alternative form of the `MQQueue` constructor has two parameters, as shown in the following example:

```
MQQueue q2 = new MQQueue("QM2", "Q2");
```

The `MQQueue` object created by this statement represents an IBM MQ queue called Q2 that is owned by queue manager QM2. The queue manager identified in this way can be the local queue manager or a remote queue manager. If it is a remote queue manager, IBM MQ must be configured so that, when the application sends a message to this destination, WebSphere MQ can route the message from the local queue manager to the remote queue manager.

The `MQQueue` constructor can also accept a queue uniform resource identifier (URI) as a single parameter. A queue URI is a string that specifies the name of an IBM MQ queue and, optionally, the name of the queue manager that owns the queue, and one or more properties of the `MQQueue` object. The following statement contains an example of a queue URI:

```
MQQueue q3 = new MQQueue("queue://QM3/Q3?persistence=2&priority=5");
```

The MQQueue object created by this statement represents an IBM MQ queue called Q3 that is owned by queue manager QM3, and all messages sent to this destination are persistent and have a priority of 5. For more information about queue URIs, see “Uniform resource identifiers (URIs)” on page 946. For an alternative way of setting the properties of an MQQueue object, see “Setting the properties of destinations.”

To create a Topic object, an application can use the MQTopic constructor, as shown in the following example:

```
MQTopic t1 = new MQTopic("Sport/Football/Results");
```

This statement creates an MQTopic object with the default values for all its properties. The object represents a topic called Sport/Football/Results.

The MQTopic constructor can also accept a topic URI as a parameter. A topic URI is a string that specifies the name of a topic and, optionally, one or more properties of the MQTopic object. The following statement contains an example of a topic URI:

```
MQTopic t2 = new MQTopic("topic://Sport/Tennis/Results?persistence=1&priority=0");
```

The MQTopic object created by this statement represents a topic called Sport/Tennis/Results, and all messages sent to this destination are nonpersistent and have a priority of 0. For more information about topic URIs, see “Uniform resource identifiers (URIs)” on page 946. For an alternative way of setting the properties of an MQTopic object, see “Setting the properties of destinations.”

### Setting the properties of destinations

An application can set the properties of a destination by calling the appropriate methods of the destination. The destination can either be an administered object or an object created dynamically at run time.

Consider the following code, for example:

```
MQQueue q1 = new MQQueue("Q1");
.
q1.setPersistence(WMQConstants.WMQ_PER_PER);
q1.setPriority(5);
```

This code creates an MQQueue object and then sets two properties of the object. The effect of setting these properties is that all messages sent to the destination are persistent and have a priority of 5.

An application can set the properties of MQTopic object in a similar way, as shown in the following example:

```
MQTopic t1 = new MQTopic("Sport/Football/Results");
.
t1.setPersistence(WMQConstants.WMQ_PER_NON);
t1.setPriority(0);
```

This code creates an MQTopic object and then sets two properties of the object. The effect of setting these properties is that all messages sent to the destination are nonpersistent and have a priority of 0.

For more information about the properties of a destination, and the methods used to set its properties, see Properties of IBM MQ classes for JMS objects.

## Building a connection in a JMS application:

To build a connection, a JMS application uses a `ConnectionFactory` object to create a `Connection` object and then starts the connection.

To create a `Connection` object, an application uses the `createConnection()` method of a `ConnectionFactory` object, as shown in the following example:

```
ConnectionFactory factory;
Connection connection;
.
.
.
connection = factory.createConnection();
```

When a JMS connection is created, the IBM MQ classes for JMS creates a connection handle (Hconn) and starts a conversation with the queue manager.

The `QueueConnectionFactory` interface and the `TopicConnectionFactory` interface each inherits the `createConnection()` method from the `ConnectionFactory` interface. You can therefore use the `createConnection()` method to create a domain specific object, as shown in the following example:

```
QueueConnectionFactory qcf;
Connection connection;
.
.
.
connection = qcf.createConnection();
```

This fragment of code creates a `QueueConnection` object. An application can now perform a domain independent operation on this object, or an operation that is applicable only to the point-to-point domain. However, if the application attempts to perform an operation that is applicable only to the publish/subscribe domain, an `IllegalStateException` exception is thrown with the following message:

```
JMSMQ1112: Operation for a domain specific object was not valid.
 Operation createProducer() is not valid for type com.ibm.mq.jms.MQTopic
```

This is because the connection was created from a domain specific connection factory.

**Note:** Note that the application process ID is used as the default user identity to be passed to the queue manager. If the application is running in client transport mode then this process ID must exist, with the relevant authorizations, on the server. If you want a different identity to be used, then use the `createConnection(username, password)` method.

The JMS specification states that a connection is created in the stopped state. Until a connection starts, a message consumer that is associated with the connection cannot receive any messages. To start a connection, an application uses the `start()` method of a `Connection` object, as shown in the following example:

```
connection.start();
```

## Creating a session in a JMS application:

To create a session, a JMS application uses the `createSession()` method of a `Connection` object.

The `createSession()` method has two parameters:

1. A parameter that specifies whether the session is transacted or not transacted
2. A parameter that specifies the acknowledgment mode for the session

For example, the following code creates a session that is not transacted and has an acknowledgment mode of `AUTO_ACKNOWLEDGE`:

```
Session session;
.
boolean transacted = false;
session = connection.createSession(transacted, Session.AUTO_ACKNOWLEDGE);
```

When a JMS session is created, the IBM MQ classes for JMS creates a connection handle (`Hconn`) and starts a conversation with the queue manager.

A `Session` object, and any `MessageProducer` or `MessageConsumer` object created from it, cannot be used concurrently by different threads of a multithreaded application. The simplest way of ensuring that these objects are not used concurrently is to create a separate `Session` object for each thread.

### *Transacted sessions in JMS applications:*

JMS applications can run local transactions by first creating a transacted session. An application can commit or roll back a transaction.

JMS applications can run local transactions. A local transaction is a transaction that involves changes only to the resources of the queue manager to which the application is connected. To run local transactions, an application must first create a transacted session by calling the `createSession()` method of a `Connection` object, specifying as a parameter that the session is transacted. Subsequently, all messages sent and received within the session are grouped into a sequence of transactions. A transaction ends when the application commits or rolls back the messages it has sent and received since the transaction began.

To commit a transaction, an application calls the `commit()` method of the `Session` object. When a transaction is committed, all messages sent within the transaction become available for delivery to other applications, and all messages received within the transaction are acknowledged so that the messaging server does not attempt to deliver them to the application again. In the point-to-point domain, the messaging server also removes the received messages from their queues.

To roll back a transaction, an application calls the `rollback()` method of the `Session` object. When a transaction is rolled back, all messages sent within the transaction are discarded by the messaging server, and all messages received within the transaction become available for delivery again. In the point-to-point domain, the messages that were received are put back on their queues and become visible to other applications again.

A new transaction starts automatically when an application creates a transacted session or calls the `commit()` or `rollback()` method. Therefore, a transacted session always has an active transaction.

When an application closes a transacted session, an implicit rollback occurs. When an application closes a connection, an implicit rollback occurs for all the connection's transacted sessions.

If an application ends without closing a connection, an implicit rollback also occurs for all the connection's transacted sessions.

A transaction is wholly contained within a transacted session. A transaction cannot span sessions. This means that it is not possible for an application to send and receive messages in two or more transacted sessions and then commit or roll back all these actions as a single transaction.

#### *Acknowledgment modes of JMS sessions:*

Every session that is not transacted has an acknowledgment mode that determines how messages received by the application are acknowledged. Three acknowledgment modes are available, and the choice of acknowledgment mode affects the design of the application.

If a session is not transacted, the way that messages received by the application are acknowledged is determined by the acknowledgment mode of the session. The three acknowledgment modes are described in the following paragraphs:

#### **AUTO\_ACKNOWLEDGE**

The session automatically acknowledges each message received by the application.

If messages are delivered synchronously to the application, the session acknowledges receipt of a message every time a Receive call completes successfully. If messages are delivered asynchronously, the session acknowledges receipt of a message every time a call to the `onMessage()` method of a message listener completes successfully.

If the application receives a message successfully, but a failure prevents acknowledgment from occurring, the message becomes available for delivery again. The application must therefore be able to handle a message that is re-delivered.

#### **DUPS\_OK\_ACKNOWLEDGE**

The session acknowledges the messages received by the application at times it selects.

Using this acknowledgment mode reduces the amount of work the session must do, but a failure that prevents message acknowledgment might result in more than one message becoming available for delivery again. The application must therefore be able to handle messages that are re-delivered.

**Restriction:** In `AUTO_ACKNOWLEDGE` and `DUPS_OK_ACKNOWLEDGE` modes, JMS does not support an application throwing an unhandled exception in a message listener. This means that messages are always acknowledged when the message listener returns, regardless of whether it was processed successfully (provided any failures are non-fatal and do not prevent the application from continuing). If you require finer control of message acknowledgment, use the `CLIENT_ACKNOWLEDGE` or transacted modes, which give the application full control of the acknowledgment functions.

#### **CLIENT\_ACKNOWLEDGE**

The application acknowledges the messages it receives by calling the `Acknowledge` method of the `Message` class.

The application can acknowledge the receipt of each message individually, or it can receive a batch of messages and call the `Acknowledge` method only for the last message it receives. When the `Acknowledge` method is called all messages received since the last time the method was called are acknowledged.

In conjunction with any of these acknowledgment modes, an application can stop and restart the delivery of messages in a session by calling the `Recover` method of the `Session` class. Messages received but previously unacknowledged are re-delivered. However, they might not be delivered in the same sequence in which they were previously delivered. In the meantime, higher priority messages might have arrived, and some of the original messages might have expired. In the point-to-point domain, some of the original messages might have been consumed by another application.



An application can determine whether a message is being re-delivered by examining the contents of the `JMSRedelivered` header field of the message. The application does this by calling the `getJMSRedelivered()` method of the `Message` class.

### Creating destinations in a JMS application:

Instead of retrieving destinations as administered objects from a Java Naming and Directory Interface (JNDI) namespace, a JMS application can use a session to create destinations dynamically at run time. An application can use a uniform resource identifier (URI) to identify an IBM MQ queue or a topic and, optionally, to specify one or more properties of a Queue or Topic object.

### Using a session to create Queue objects

To create a Queue object, an application can use the `createQueue()` method of a Session object, as shown in the following example:

```
Session session;
.
Queue q1 = session.createQueue("Q1");
```

This code creates a Queue object with the default values for all its properties. The object represents an IBM MQ queue called Q1 that belongs to the local queue manager. This queue can be a local queue, an alias queue, or a remote queue definition.

The `createQueue()` method also accepts a queue URI as a parameter. A queue URI is a string that specifies the name of an IBM MQ queue and, optionally, the name of the queue manager that owns the queue and one or more properties of the Queue object. The following statement contains an example of a queue URI:

```
Queue q2 = session.createQueue("queue://QM2/Q2?persistence=2&priority=5");
```

The Queue object created by this statement represents an IBM MQ queue called Q2 that is owned by a queue manager called QM2, and all messages sent to this destination are persistent and have a priority of 5. The queue manager identified in this way can be the local queue manager or a remote queue manager. If it is a remote queue manager, IBM MQ must be configured so that, when the application sends a message to this destination, WebSphere MQ can route the message from the local queue manager to queue manager QM2. For more information about URIs, see “Uniform resource identifiers (URIs)” on page 946.

Note that the parameter on the `createQueue()` method contains provider specific information. Therefore, using the `createQueue()` method to create a Queue object, instead of retrieving a Queue object as an administered object from a JNDI namespace, might make your application less portable.

An application can create a `TemporaryQueue` object by using the `createTemporaryQueue()` method of a Session object, as shown in the following example:

```
TemporaryQueue q3 = session.createTemporaryQueue();
```

Although a session is used to create a temporary queue, the scope of a temporary queue is the connection that was used to create the session. Any of the connection's sessions can create message producers and message consumers for the temporary queue. The temporary queue remains until the connection ends or the application explicitly deletes the temporary queue by using the `TemporaryQueue.delete()` method, whichever is the sooner.

When an application creates a temporary queue, IBM MQ classes for JMS creates a dynamic queue in the queue manager to which the application is connected. The `TEMPMODEL` property of the connection factory specifies the name of the model queue that is used to create the dynamic queue, and the `TEMPQPREFIX` property of the connection factory specifies the prefix that is used to form the name of the dynamic queue.

## Using a session to create Topic objects

To create a Topic object, an application can use the createTopic() method of a Session object, as shown in the following example:

```
Session session;
.
Topic t1 = session.createTopic("Sport/Football/Results");
```

This code creates an Topic object with the default values for all its properties. The object represents a topic called Sport/Football/Results.

The createTopic() method also accepts a topic URI as a parameter. A topic URI is a string that specifies the name of a topic and, optionally, one or more properties of the Topic object. The following code contains an example of a topic URI:

```
String uri = "topic://Sport/Tennis/Results?persistence=1&priority=0";
Topic t2 = session.createTopic(uri);
```

The Topic object created by this code represents a topic called Sport/Tennis/Results, and all messages sent to this destination are nonpersistent and have a priority of 0. For more information about topic URIs, see "Uniform resource identifiers (URIs)."

Note that the parameter on the createTopic() method contains provider specific information. Therefore, using the createTopic() method to create a Topic object, instead of retrieving a Topic object as an administered object from a JNDI namespace, might make your application less portable.

An application can create a TemporaryTopic object by using the createTemporaryTopic() method of a Session object, as shown in the following example:

```
TemporaryTopic t3 = session.createTemporaryTopic();
```

Although a session is used to create a temporary topic, the scope of a temporary topic is the connection that was used to create the session. Any of the connection's sessions can create message producers and message consumers for the temporary topic. The temporary topic remains until the connection ends or the application explicitly deletes the temporary topic by using the TemporaryTopic.delete() method, whichever is the sooner.

When an application creates a temporary topic, IBM MQ classes for JMS creates a topic with a name that commences with the characters TEMP/*tempTopicPrefix*, where *tempTopicPrefix* is the value of the TEMPTOPICPREFIX property of the connection factory.

## Uniform resource identifiers (URIs)

A queue URI is a string that specifies the name of an IBM MQ queue and, optionally, the name of the queue manager that owns the queue and one or more properties of the Queue object created by the application. A topic URI is a string that specifies the name of a topic and, optionally, one or more properties of the Topic object created by the application.

A queue URI has the following format:

```
queue://[qMgrName]/qName [? propertyName1 = propertyValue1
& propertyName2 = propertyValue2
&...]
```

A topic URI has the following format:

```
topic://topicName [? propertyName1 = propertyValue1
& propertyName2 = propertyValue2
&...]
```

The variables in these formats have the following meanings:

***qMgrName***

The name of the queue manager that owns the queue identified by the URI.

The queue manager can be the local queue manager or a remote queue manager. If it is a remote queue manager, IBM MQ must be configured so that, when an application sends a message to the queue, WebSphere MQ can route the message from the local queue manager to the remote queue manager.

If no name is specified, the local queue manager is assumed.

***qName***

The name of the IBM MQ queue.

The queue can be a local queue, an alias queue, or a remote queue definition.

For the rules for creating queue names, see Rules for naming IBM MQ objects.

***topicName***

The name of the topic.

For the rules for creating topic names, see Rules for naming IBM MQ objects. Avoid the use of the wildcard characters +, #, \*, and ? in topic names. Topic names containing these characters can cause unexpected results when you subscribe to them. See Using topic strings.

***propertyName1, propertyName2, ...***

The names of the properties of the Queue or Topic object created by the application. Table 133 lists the valid property names that can be used in a URI.

If no properties are specified, the Queue or Topic object has the default values for all its properties.

***propertyValue1, propertyValue2, ...***

The values of the properties of the Queue or Topic object created by the application. Table 133 lists the valid property values that can be used in a URI.

Brackets ([]) denotes an optional component, and the ellipsis (...) means that the list of property name-value pairs, if present, can contain one or more name-value pairs.

Table 133 lists the valid property names and valid values that can be used in queue and topic URIs. Although the IBM MQ JMS administration tool uses symbolic constants for the values of properties, URIs cannot contain symbolic constants.

*Table 133. Property names and valid values for use in queue and topic URIs*

Property name	Description	Valid values
CCSID	How the character data in the body of a message is represented when IBM MQ classes for JMS forwards the message to the destination	<ul style="list-style-type: none"> <li>Any coded character set identifier supported by IBM MQ.</li> </ul>
encoding	How the numeric data in the body of a message is represented when IBM MQ classes for JMS forwards the message to the destination	<ul style="list-style-type: none"> <li>Any valid value for the <i>Encoding</i> field in an IBM MQ message descriptor.</li> </ul>
expiry	The time to live for messages sent to the destination	<ul style="list-style-type: none"> <li>-2 - As specified on the send() call or, if not specified on the send() call, the default time to live of the message producer.</li> <li>0 - A message sent to the destination never expires.</li> <li>A positive integer specifying the time to live in milliseconds.</li> </ul>

Table 133. Property names and valid values for use in queue and topic URIs (continued)

Property name	Description	Valid values
multicast	The multicast setting for a topic when using a real-time connection to a broker	The following list contains the valid values. Associated with each value is the corresponding value of the MULTICAST property as used in the IBM MQ JMS administration tool. For a description of the MULTICAST property and its valid values, see Properties of IBM MQ classes for JMS objects. <ul style="list-style-type: none"> <li>• -1 - ASCF</li> <li>• 0 - DISABLED</li> <li>• 3 - NOTR</li> <li>• 5 - RELIABLE</li> <li>• 7 - ENABLED</li> </ul>
persistence	The persistence of messages sent to the destination	<ul style="list-style-type: none"> <li>• -2 - As specified on the send() call or, if not specified on the send() call, the default persistence of the message producer.</li> <li>• -1 - As specified by the <i>DefPersistence</i> attribute of the IBM MQ queue or topic.</li> <li>• 1 - Nonpersistent.</li> <li>• 2 - Persistent.</li> <li>• 3 - Equivalent to the value HIGH for the PERSISTENCE property as used in the IBM MQ JMS administration tool. For an explanation of this value, see “JMS persistent messages” on page 970.</li> </ul>
priority	The priority of messages sent to the destination	<ul style="list-style-type: none"> <li>• -2 - As specified on the send() call or, if not specified on the send() call, the default priority of the message producer.</li> <li>• -1 - As specified by the <i>DefPriority</i> attribute of the IBM MQ queue or topic.</li> <li>• An integer in the range 0-9 specifying the priority of messages sent to the destination.</li> </ul>
targetClient	Whether messages sent to the destination contain an MQRFH2 header	<ul style="list-style-type: none"> <li>• 0 - Messages contain an MQRFH2 header.</li> <li>• 1 - Messages do not contain an MQRFH2 header.</li> </ul>

For example, the following URI identifies an IBM MQ queue called Q1 that is owned by the local queue manager. A Queue object created using this URI has the default values for all its properties.

```
queue:///Q1
```

The following URI identifies an IBM MQ queue called Q2 that is owned by a queue manager called QM2. All messages sent to this destination have a priority of 6. The remaining properties of the Queue object created using this URI have their default values.

```
queue://QM2/Q2?priority=6
```

The following URI identifies a topic called Sport/Athletics/Results. All messages sent to this destination are nonpersistent and have a priority of 0. The remaining properties of the Topic object created using this URI have their default values.

```
topic://Sport/Athletics/Results?persistence=1&priority=0
```

## Sending messages in a JMS application:

Before a JMS application can send messages to a destination, it must first create a `MessageProducer` object for the destination. To send a message to the destination, the application creates a `Message` object and then calls the `send()` method of the `MessageProducer` object.

An application uses a `MessageProducer` object to send messages. An application normally creates a `MessageProducer` object for a specific destination, which can be a queue or a topic, so that all messages sent using the message producer are sent to the same destination. Therefore, before an application can create a `MessageProducer` object, it must first create a `Queue` or `Topic` object. For information about how to create a `Queue` or `Topic` object, see the following topics:

- “Using JNDI to retrieve administered objects in a JMS application” on page 931
- “Using the IBM JMS extensions” on page 932
- “Using the IBM MQ JMS extensions” on page 939
- “Creating destinations in a JMS application” on page 945

To create a `MessageProducer` object, an application uses the `createProducer()` method of a `Session` object, as shown in the following example:

```
MessageProducer producer = session.createProducer(destination);
```

The parameter `destination` is a `Queue` or `Topic` object that the application has created previously.

Before an application can send a message, it must create a `Message` object. The body of a message contains the application data, and JMS defines five types of message body:

- Bytes
- Map
- Object
- Stream
- Text

Each type of message body has its own JMS interface, which is a sub-interface of the `Message` interface, and a method in the `Session` interface for creating a message with that type of body. For example, the interface for a text message is called `TextMessage`, and an application uses the `createTextMessage()` method of a `Session` object to create a text message, as shown in the following statement:

```
TextMessage outMessage = session.createTextMessage(outString);
```

For more information about messages and message bodies, see “JMS messages” on page 875.

To send a message, an application uses the `send()` method of a `MessageProducer` object, as shown in the following example:

```
producer.send(outMessage);
```

An application can use the `send()` method to send messages in either messaging domain. The nature of the destination determines which messaging domain is used. However, `TopicPublisher`, the sub-interface of `MessageProducer` that is specific to the publish/subscribe domain, also has a `publish()` method, which can be used instead of the `send()` method. The two methods are functionally the same.

An application can create a `MessageProducer` object with no specified destination. In this case, the application must specify the destination when calling the `send()` method.

If an application sends a message within a transaction, the message is not delivered to its destination until the transaction is committed. This means that an application cannot send a message and receive a reply to the message within the same transaction.

A destination can be configured so that when an application sends messages to it, IBM MQ classes for JMS forwards the message and returns control back to the application without determining whether the queue manager has received the message safely. This is sometimes referred to as *asynchronous put*. For more information, see “Putting messages asynchronously in IBM MQ classes for JMS” on page 993.

### Receiving messages in a JMS application:

An application uses a message consumer to receive messages. A durable topic subscriber is a message consumer that receives all messages sent to a destination, including those sent while the consumer is inactive. An application can select which messages it wants to receive by using a message selector, and can receive messages asynchronously by using a message listener.

An application uses a `MessageConsumer` object to receive messages. An application creates a `MessageConsumer` object for a specific destination, which can be a queue or a topic, so that all messages received using the message consumer are received from the same destination. Therefore, before an application can create a `MessageConsumer` object, it must first create a `Queue` or `Topic` object. For information about how to create a `Queue` or `Topic` object, see the following topics:

- “Using JNDI to retrieve administered objects in a JMS application” on page 931
- “Using the IBM JMS extensions” on page 932
- “Using the IBM MQ JMS extensions” on page 939
- “Creating destinations in a JMS application” on page 945

To create a `MessageConsumer` object, an application uses the `createConsumer()` method of a `Session` object, as shown in the following example:

```
MessageConsumer consumer = session.createConsumer(destination);
```

The parameter `destination` is a `Queue` or `Topic` object that the application has created previously.

The application then uses the `receive()` method of the `MessageConsumer` object to receive a message from the destination, as shown in the following example:

```
Message inMessage = consumer.receive(1000);
```

The parameter on the `receive()` call specifies how long in milliseconds the method waits for a suitable message to arrive if no message is available immediately. If you omit this parameter, the call blocks indefinitely until a suitable message arrives. If you do not want the application to wait for a message, use the `receiveNoWait()` method instead.

The `receive()` method returns a message of a specific type. For example, when an application receives a text message, the object returned by the `receive()` call is a `TextMessage` object.

However, the declared type of object returned by a `receive()` call is a `Message` object. Therefore, in order to extract the data from the body of a message that has just been received, the application must cast from the `Message` class to the more specific subclass, such as `TextMessage`. If the type of the message is not known, the application can use the `instanceof` operator to determine the type. It is always good practice for an application to determine the type of a message before casting so that errors can be handled gracefully.

The following code uses the `instanceof` operator and shows how to extract the data from the body of a text message:

```
if (inMessage instanceof TextMessage) {
 String replyString = ((TextMessage) inMessage).getText();
 .
 .
 .
}
```

```

} else {
 // Print error message if Message was not a TextMessage.
 System.out.println("Reply message was not a TextMessage");
}

```

If an application sends a message within a transaction, the message is not delivered to its destination until the transaction is committed. This means that an application cannot send a message and receive a reply to the message within the same transaction.

If a message consumer receives messages from a destination that is configured for read ahead, any nonpersistent messages that are in the read ahead buffer when the application ends are discarded.

In the publish/subscribe domain, JMS identifies two types of message consumer, nondurable topic subscriber and durable topic subscriber, which are described in the following two sections.

### **Nondurable topic subscribers**

A nondurable topic subscriber receives only those messages that are published while the subscriber is active. A nondurable subscription starts when an application creates a nondurable topic subscriber and ends when the application closes the subscriber, or when the subscriber falls out of scope. As an extension in IBM MQ classes for JMS, a nondurable topic subscriber also receives retained publications.

To create a nondurable topic subscriber, an application can use the domain independent `createConsumer()` method, specifying a `Topic` object as the destination. Alternatively, an application can use the domain specific `createSubscriber()` method, as shown in the following example:

```
TopicSubscriber subscriber = session.createSubscriber(topic);
```

The parameter `topic` is a `Topic` object that the application has created previously.

### **Durable topic subscribers**

**Restriction:** An application cannot create durable topic subscribers when using a real-time connection to a broker.

A durable topic subscriber receives all messages that are published during the life of a durable subscription. These messages include all those that are published while the subscriber is not active. As an extension in IBM MQ classes for JMS, a durable topic subscriber also receives retained publications.

To create a durable topic subscriber, an application uses the `createDurableSubscriber()` method of a `Session` object, as shown in the following example:

```
TopicSubscriber subscriber = session.createDurableSubscriber(topic, "D_SUB_000001");
```

On the `createDurableSubscriber()` call, the first parameter is a `Topic` object that the application has created previously, and the second parameter is a name that is used to identify the durable subscription.

The session used to create a durable topic subscriber must have an associated client identifier. The client identifier associated with a session is the same as the client identifier for the connection that is used to create the session. The client identifier can be specified by setting the `CLIENTID` property of the `ConnectionFactory` object. Alternatively, an application can specify the client identifier by calling the `setClientID()` method of the `Connection` object.

The name that is used to identify a durable subscription must be unique only within the client identifier, and therefore the client identifier forms part of the full, unique identifier of a durable subscription. To continue using a durable subscription that was created previously, an application must create a durable topic subscriber using a session with the same client identifier as that associated with the durable subscription, and using the same subscription name.

A durable subscription starts when an application creates a durable topic subscriber using a client identifier and subscription name for which no durable subscription currently exists. However, a durable subscription does not end when the application closes the durable topic subscriber. To end a durable subscription, an application must call the `unsubscribe()` method of a `Session` object that has the same client identifier as that associated with the durable subscription. The parameter on the `unsubscribe()` call is the subscription name, as shown in the following example:

```
session.unsubscribe("D_SUB_000001");
```

The scope of a durable subscription is a queue manager. If a durable subscription exists on one queue manager, and an application connected to another queue manager creates a durable subscription with the same client identifier and subscription name, the two durable subscriptions are completely independent.

### Message selectors

An application can specify that only those messages that satisfy certain criteria are returned by successive `receive()` calls. When creating a `MessageConsumer` object, the application can specify a Structured Query Language (SQL) expression that determines which messages are retrieved. This SQL expression is called a *message selector*. The message selector can contain the names of JMS message header fields and message properties. For information about how to construct a message selector, see “Message selectors in JMS” on page 875.

The following example shows how an application can select messages based on a user defined property called `myProp`:

```
MessageConsumer consumer;
.
consumer = session.createConsumer(destination, "myProp = 'blue'");
```

The JMS specification does not allow an application to change the message selector of a message consumer. After an application creates a message consumer with a message selector, the message selector remains for the life of that consumer. If an application requires more than one message selector, the application must create a message consumer for each message selector.

Note that, when an application is connected to a Version 7 queue manager, the `MSGSELECTION` property of the connection factory has no effect. To optimize performance, all message selection is done by the queue manager.

### Suppressing local publications

An application can create a message consumer that ignores publications published on the consumer's own connection. The application does this by setting the third parameter on a `createConsumer()` call to `true`, as shown in the following example:

```
MessageConsumer consumer = session.createConsumer(topic, null, true);
```

On a `createDurableSubscriber()` call, the application does this by setting the fourth parameter to `true`, as shown in the following example

```
String selector = "company = 'IBM'";
TopicSubscriber subscriber = session.createDurableSubscriber(topic, "D_SUB_000001",
 selector, true);
```

### Asynchronous delivery of messages

An application can receive messages asynchronously by registering a message listener with a message consumer. The message listener has a method called `onMessage`, which is called asynchronously when a suitable message is available and whose purpose is to process the message. The following code illustrates the mechanism:



```

import javax.jms.*;

public class MyClass implements MessageListener
{
 // The method that is called asynchronously when a suitable message is available
 public void onMessage(Message message)
 {
 System.out.println("Message is "+message);

 // The code to process the message
 .
 .
 .
 }
}
.
.
.
// Main program (possibly in another class)
.
// Creating the message listener
MyClass listener = new MyClass();

// Registering the message listener with a message consumer
consumer.setMessageListener(listener);

// The main program now continues with other processing

```

An application can use a session either for receiving messages synchronously using `receive()` calls, or for receiving messages asynchronously using message listeners, but not for both. If an application needs to receive messages synchronously and asynchronously, it must create separate sessions.

Once a session is set up to receive messages asynchronously, the following methods cannot be called on that session or on objects created from that session:

- `MessageConsumer.receive()`
- `MessageConsumer.receive(long)`
- `MessageConsumer.receiveNoWait()`
- `Session.acknowledge()`
- `MessageProducer.send(Destination, Message)`
- `MessageProducer.send(Destination, Message, int, int, long)`
- `MessageProducer.send(Message)`
- `MessageProducer.send(Message, int, int, long)`
- `MessageProducer.send(Destination, Message, CompletionListener)`
- `MessageProducer.send(Destination, Message, int, int, long, CompletionListener)`
- `MessageProducer.send(Message, CompletionListener)`
- `MessageProducer.send(Message, int, int, long, CompletionListener)`
- `Session.commit()`
- `Session.createBrowser(Queue)`
- `Session.createBrowser(Queue, String)`
- `Session.createBytesMessage()`
- `Session.createConsumer(Destination)`
- `Session.createConsumer(Destination, String, boolean)`
- `Session.createDurableSubscriber(Topic, String)`
- `Session.createDurableSubscriber(Topic, String, String, boolean)`
- `Session.createMapMessage()`

- Session.createMessage()
- Session.createObjectMessage()
- Session.createObjectMessage(Serializable)
- Session.createProducer(Destination)
- Session.createQueue(String)
- Session.createStreamMessage()
- Session.createTemporaryQueue()
- Session.createTemporaryTopic()
- Session.createTextMessage()
- Session.createTextMessage(String)
- Session.createTopic()
- Session.getAcknowledgeMode()
- Session.getMessageListener()
- Session.getTransacted()
- Session.rollback()
- Session.unsubscribe(String)

If any of these methods are called, a `JMSEException` containing the message:

JMSCC0033: A synchronous method call is not permitted when a session is being used asynchronously: 'method name' is thrown.

### Receiving poison messages

An application can receive a message that cannot be processed. There can be several reasons why the message cannot be processed, for example the message might have an incorrect format. Such messages are described as poison messages and require special handling to prevent the message being recursively processed.

For details on how to handle poison messages, see “Handling poison messages in IBM MQ classes for JMS” on page 956.

### Retrieval of subscription user data:

If the messages that an IBM MQ classes for JMS application is consuming from a queue are put by an administratively defined durable subscription, the application needs to access the user data information that is associated with the subscription. This information is added to the message as a property.

For Continuous Delivery, from IBM MQ Version 9.0.2, when a message is consumed from a queue that contains an RFH2 header with the MQPS folder, the value that is associated with the `Sud` key, if it exists, is added as a `String` property to the JMS Message object returned to the IBM MQ classes for JMS application. To enable the retrieval of this property from the message, the constant `JMS_IBM_SUBSCRIPTION_USER_DATA` in the `JmsConstants` interface can be used with the method `javax.jms.Message.getStringProperty(java.lang.String)` to get the subscription user data.

From IBM MQ Version 9.0.0, Fix Pack 2 the constant `JMS_IBM_SUBSCRIPTION_USER_DATA` is also available in Long Term Support.

In the following example, an administrative durable subscription is defined by using the MQSC command **DEFINE SUB**:

```
DEFINE SUB('MY.SUBSCRIPTION') TOPICSTR('PUBLIC') DEST('MY.SUBSCRIPTION.Q')
USERDATA('Administrative durable subscription to put message to the queue MY.SUBSCRIPTION.Q')
```

Copies of messages that are published to the topic string PUBLIC are put to the queue, MY.SUBSCRIPTION.Q. The user data that is associated with the durable subscription is then added as a property to the message, which is stored in the MQPS folder of the RFH2 header with the key Sud.

The IBM MQ classes for JMS application can call:

```
javax.jms.Message.getStringProperty(JmsConstants.JMS_IBM_SUBSCRIPTION_USER_DATA);
```

The following String is then returned:

```
Administrative durable subscription to put message to the queue MY.SUBSCRIPTION.Q
```

.

#### **Related concepts:**

“The MQRFH2 header and JMS” on page 880

This collection of topics describes the MQRFH Version 2 header, which carries JMS-specific data that is associated with the message content. The MQRFH Version 2 header is extensible, and can also carry additional information that is not directly associated with JMS. However, this section covers only its use by JMS.

#### **Related information:**

Defining an administrative subscription

DEFINE SUB

Interface JmsConstants

#### **Closing down an IBM MQ classes for JMS application:**

It is important for an IBM MQ classes for JMS application to close certain JMS objects explicitly before stopping. Finalizers might not be called, so do not rely on them to free resources. Do not allow an application to terminate with compressed trace active.

Garbage collection alone cannot release all IBM MQ classes for JMS and IBM MQ resources in a timely manner, especially if an application creates many short lived JMS objects at the session level or lower. It is therefore important for an application to close a Connection, Session, MessageConsumer, or MessageProducer object when it is no longer required.

If an application ends without closing a Connection, an implicit rollback occurs for all the connection's transacted sessions. To ensure any changes made by the application are committed, close the Connection explicitly before closing the application.

Do not use finalizers in an application to close JMS objects. Because finalizers might not be called, resources might not be freed. When a Connection is closed it closes all the Sessions that were created from it. Similarly, the MessageConsumers and MessageProducers created from a Session are closed when the Session is closed. However, consider closing Sessions, MessageConsumers, and MessageProducers explicitly to ensure resources are freed in a timely manner.

If trace compression is activated, System.Halt() shutdowns and abnormal, uncontrolled JVM terminations are likely to result in a corrupt trace file. Where possible, turn off the trace facility when you have collected the trace information you need. If you are tracing an application up to an abnormal end, use uncompressed trace output.

**Note:** To disconnect from a queue manager, a JMS application invokes the close() method on the connection object.

## Handling poison messages in IBM MQ classes for JMS:

A poison message is one which cannot be processed by a receiving MDB application. If a poison message is encountered, the JMS MessageConsumer and ConnectionConsumer objects can requeue it according to two queue properties, BOQUEUE, and BOTHRESH.

Sometimes, a badly formatted message arrives on a queue. In this context, badly formatted means that the receiving application cannot process the message correctly. Such a message can cause the receiving application to fail and to back out this badly formatted message. The message can then be repeatedly delivered to the input queue and repeatedly backed out by the application. These messages are known as *poison messages*. The JMS MessageConsumer object detects poison messages and reroutes them to an alternative destination.

The IBM MQ queue manager keeps a record of the number of times that each message has been backed out. When this number reaches a configurable threshold value, the message consumer requeues the message to a named backout queue. If this requeuing fails for any reason, the message is removed from the input queue and either requeued to the dead-letter queue, or discarded. See "Removing messages from the queue in ASF" on page 1005 for more details.

There is a difference between the way in which poison messages are requeued by MessageConsumers and ConnectionConsumers. ConnectionConsumers are able to requeue poison messages without affecting message delivery. The requeue process takes place outside of any unit of work associated with actual message delivery to application code. This is possible because of the multi-threaded nature of ConnectionConsumer operation.

MessageConsumers, however, are single threaded below the Session level, and any requeuing of poison messages takes place within the current unit of work. This does not affect the operation of the application, however when poison messages are requeued under a transacted or Client\_acknowledge Session, the requeue action itself is not committed until the current unit of work is committed by the application code or, if appropriate, the application container code.

JMS ConnectionConsumer objects handle poison messages in the same way and using the same queue properties. If multiple connection consumers are monitoring the same queue, it is possible that the poison message may be delivered to an application more times than the threshold value before the requeue occurs. This behavior is due to the way individual connection consumers monitor queues and requeue poison messages.

The threshold value and the name of the back out queue are attributes of an IBM MQ queue. The names of the attributes are BackoutThreshold and BackoutRequeueQName. The queue they apply to is as follows:

- For point-to-point messaging, this is the underlying local queue. This is important when message consumers and connection consumers use queue aliases.
- For publish/subscribe messaging in IBM MQ messaging provider normal mode, the Topic's managed queue is created from the model queue.
- For publish/subscribe messaging in IBM MQ messaging provider migration mode, it is the CCSUB queue defined on the TopicConnectionFactory object, or the CCDSUB queue defined on the Topic object.

IBM MQ classes for JMS queries the BackoutThreshold and BackoutRequeueQName of the queue. You must therefore grant inquire access on the queue to the user running the application. If the target queue is a cluster queue, grant inquire, browse and get access.

To set the BackoutThreshold and BackoutRequeueQName attributes, issue the following MQSC command:

```
ALTER QLOCAL(your.queue.name) BOTHRESH(threshold value) BOQUEUE(your.backout.queue.name)
```

If the `BackoutThreshold` attribute is set to a value other than zero, to avoid unexpected behavior set the `BackoutRequeueQName` attribute to a valid queue name.

For publish/subscribe messaging, if your system creates a dynamic queue for each subscription, these attribute values are obtained from the IBM MQ classes for JMS model queue, `SYSTEM.JMS.MODEL.QUEUE`. To alter these settings, use:

```
ALTER QMODEL(SYSTEM.JMS.MODEL.QUEUE) BOTHRESH(threshold value) BOQUEUE(your.backout.queue.name)
```

If the backout threshold value is zero, poison message handling is disabled, and poison messages remain on the input queue. Otherwise, when the backout count reaches the threshold value, the message is sent to the named backout queue. If the backout count reaches the threshold value, but the message cannot go to the backout queue, the message is sent to the dead-letter queue or it is discarded. This situation occurs if the backout queue is not defined, or if the `MessageConsumer` object cannot send the message to the backout queue. See “Removing messages from the queue in ASF” on page 1005 for further details.

#### **Related concepts:**

“Handling poison messages in ASF” on page 1006

Within the Application Server Facilities, poison message handling is handled slightly differently to elsewhere in IBM MQ classes for JMS.

#### **Logging errors in IBM MQ classes for JMS:**

Information about runtime problems that might require corrective action by the user is written to the IBM MQ classes for JMS log.

For example, if an application attempts to set a property of a connection factory, but the name of the property is not recognized, IBM MQ classes for JMS writes information about the problem to its log.

By default, the file containing the log is called `mjms.log` and is in the current working directory. However, you can change the name and location of the log file by setting the `com.ibm.msg.client.commonservices.log.outputName` property in the IBM MQ classes for JMS configuration file. For information about the IBM MQ classes for JMS configuration file, see “The IBM MQ classes for JMS configuration file” on page 837, and for more detail about valid values for the `com.ibm.msg.client.commonservices.log.outputName` property, see “Logging and IBM MQ classes for JMS” on page 863.

#### **Exceptions in IBM MQ classes for JMS:**

An IBM MQ classes for JMS application must be able to handle exceptions that are thrown by a JMS API calls or delivered to an exception handler.

IBM MQ classes for JMS reports runtime problems by throwing exceptions. `JMSEException` is the root class for exceptions thrown by JMS methods, and catching `JMSEException` exceptions provides a generic way of handling all JMS related exceptions.

Every `JMSEException` exception encapsulates the following information:

- A provider specific exception message, which an application obtains by calling the `Throwable.getMessage()` method.
- A provider specific error code, which an application obtains by calling the `JMSEException.getErrorCode()` method.
- A linked exception. An exception thrown by a JMS API call is often the result of a lower level problem, which is reported by another exception linked to this exception. An application obtains a linked exception by calling the `JMSEException.getLinkedException()` or the `Throwable.getCause()` method.

Most exceptions thrown by IBM MQ classes for JMS are instances of subclasses of `JMSEException`. These subclasses implement the `com.ibm.msg.client.jms.JmsExceptionDetail` interface, which provides the following additional information:

- An explanation of the exception message, which an application obtains by calling the `JmsExceptionDetail.getExplanation()` method.
- A recommended user response to the exception, which an application obtains by calling the `JmsExceptionDetail.getUserAction()` method.
- The keys for the message inserts in the exception message. An application obtains an iterator for all the keys by calling the `JmsExceptionDetail.getKeys()` method.
- The message inserts in the exception message. For example, a message insert might be the name of the queue that caused the exception, and it might be useful for an application to be able to access that name. An application obtains the message insert corresponding to a specified key by calling the `JmsExceptionDetail.getValue()` method.

All the methods in the `JmsExceptionDetail` interface might return null if no details are available.

For example, if an application tries to create a message producer for a IBM MQ queue that does not exist, an exception is thrown with the following information:

```
Message : JMSWMQ2008: Failed to open MQ queue 'Q_test'.
Class : class com.ibm.msg.client.jms.DetailedInvalidDestinationException
Error Code : JMSWMQ2008
Explanation : JMS attempted to perform an MQOPEN, but IBM MQ reported an
 error.
User Action : Use the linked exception to determine the cause of this error. Check
 that the specified queue and queue manager are defined correctly.
```

The exception thrown, `com.ibm.msg.client.jms.DetailedInvalidDestinationException`, is a subclass of `javax.jms.InvalidDestinationException` and implements the `com.ibm.msg.client.jms.JmsExceptionDetail` interface.

## Linked exceptions

A linked exception provides further information about a runtime problem. Therefore, for each `JMSEException` exception that is thrown, an application should check the linked exception. The linked exception itself might have another linked exception, and so the linked exceptions form a chain leading back to the original underlying problem. A linked exception is implemented by using the chained exception mechanism of the `java.lang.Throwable` class, and an application obtains a linked exception by calling the `Throwable.getCause()` method. For a `JMSEException` exception, the `getLinkedException()` method actually delegates to the `Throwable.getCause()` method.

For example, if an application specifies an incorrect port number when connecting to a queue manager, the exceptions form the following chain:

```
com.ibm.msg.client.jms.DetailIllegalStateException
|
+- -->
com.ibm.mq.MQException
|
+- -->
com.ibm.mq.jmqi.JmqiException
|
+- -->
java.net.ConnectionException
```

Typically, each exception in a chain is thrown from a different layer in the code. For example, the exceptions in the preceding chain are thrown by the following layers:

- The first exception, an instance of a subclass of `JMSEException`, is thrown by the common layer in IBM MQ classes for JMS.

- The next exception, an instance of `com.ibm.mq.MQException`, is thrown by the IBM MQ messaging provider.
- The next exception, an instance of `com.ibm.mq.jmqi.JmqiException`, is thrown by the common Java interface to the MQI.
- The final exception, an instance of `java.net.ConnectionException`, is thrown by the Java class library.

For more information about the layered architecture of IBM MQ classes for JMS, see [Layered architecture](#).

Using code similar to the following code, an application can iterate through this chain to extract all the appropriate information:

```
import com.ibm.msg.client.jms.JmsExceptionDetail;
import com.ibm.mq.MQException;
import com.ibm.mq.jmqi.JmqiException;
import javax.jms.JMSEException;
.
.
.
catch (JMSEException je) {
System.err.println("Caught JMSEException");

// Check for linked exceptions in JMSEException
Throwable t = je;
while (t != null) {
// Write out the message that is applicable to all exceptions
System.err.println("Exception Msg: " + t.getMessage());
// Write out the exception stack trace
t.printStackTrace(System.err);

// Add on specific information depending on the type of exception
if (t instanceof JMSEException) {
JMSEException je1 = (JMSEException) t;
System.err.println("JMS Error code: " + je1.getErrorCode());

if (t instanceof JmsExceptionDetail){
JmsExceptionDetail jed = (JmsExceptionDetail)je1;
System.err.println("JMS Explanation: " + jed.getExplanation());
System.err.println("JMS Explanation: " + jed.getUserAction());
}
} else if (t instanceof MQException) {
MQException mqe = (MQException) t;
System.err.println("WMQ Completion code: " + mqe.getCompCode());
System.err.println("WMQ Reason code: " + mqe.getReason());
} else if (t instanceof JmqiException){
JmqiException jmqie = (JmqiException)t;
System.err.println("WMQ Log Message: " + jmqie.getWmqLogMessage());
System.err.println("WMQ Explanation: " + jmqie.getWmqMsgExplanation());
System.err.println("WMQ Msg Summary: " + jmqie.getWmqMsgSummary());
System.err.println("WMQ Msg User Response: "
+ jmqie.getWmqMsgUserResponse());
System.err.println("WMQ Msg Severity: " + jmqie.getWmqMsgSeverity());
}

// Get the next cause
t = t.getCause();
}
}
```

Note that an application should always check the type of each exception in a chain because the type of exception can vary and exceptions of different types encapsulate different information.

## Obtaining IBM MQ specific information about a problem

Instances of `com.ibm.mq.MQException` and `com.ibm.mq.jmqi.JmqiException` encapsulate IBM MQ specific information about a problem.

An `MQException` exception encapsulates the following information:

- A completion code, which an application obtains by calling the `getCompCode()` method
- A reason code, which an application obtains by calling the `getReason()` method

A `JmqiException` exception also encapsulates a completion code and a reason code. Additionally, however, a `JmqiException` exception encapsulates the information in an AMQ *nnnn* or CSQ *nnnn* message, if one is associated with the exception. By calling the appropriate methods of the exception, an application can obtain the various components of this message, such as the severity, explanation, and user response.

For examples of how to use of the methods mentioned in this section, see the sample code in “Linked exceptions” on page 958.

## Upgrading from previous versions of IBM MQ classes for JMS

Compared to previous versions of IBM MQ classes for JMS, most error codes and exception messages changed in Version 7.0. The reason for these changes is that from Version 7.0, IBM MQ classes for JMS has a layered architecture and exceptions are thrown from different layers in the code.

For example, if an application tries to connect to a queue manager that does not exist, a previous version of IBM MQ classes for JMS threw a `JMSEException` exception with the following information:

```
MQJMS2005: Failed to create MQQueueManager for 'localhost:QM_test'.
```

This exception contained a linked `MQException` exception with the following information:

```
MQJE001: Completion Code 2, Reason 2058
```

By comparison in the same circumstances, Version 7.0 of IBM MQ classes for JMS throws a `JMSEException` exception with the following information:

```
Message : JMSWMQ0018: Failed to connect to queue manager 'QM_test' with
connection mode 'Client' and host name 'localhost'.
Class : class com.ibm.msg.client.jms.DetailedJMSEException
Error Code : JMSWMQ0018
Explanation : null
User Action : Check the queue manager is started and if running in client mode,
check there is a listener running. Please see the linked exception
for more information.
```

This exception contains a linked `MQException` exception with the following information:

```
Message : JMSCMQ0001: IBM MQ call failed with compcode '2' ('MQCC_FAILED')
reason '2058' ('MQRC_Q_MGR_NAME_ERROR').
Class : class com.ibm.mq.MQException
Completion Code : 2
Reason Code : 2058
```

If your application parses or tests exception messages returned by the `Throwable.getMessage()` method, or error codes returned by the `JMSEException.getErrorCode()` method, and you are upgrading from a release before Version 7.0, your application probably needs to be modified in order to use Version 7.0 or later of IBM MQ classes for JMS.



## Exception listeners

An application can register an exception listener with a Connection object. Subsequently, if a problem occurs that makes the connection unusable, IBM MQ classes for JMS delivers an exception to the exception listener by calling its onException() method. The application then has the opportunity to reestablish the connection. IBM MQ classes for JMS can also deliver an exception to the exception listener if a problem occurs while trying to deliver a message asynchronously.

From IBM MQ Version 8.0.0, Fix Pack 2, to maintain behavior for current JMS applications that configure a JMS MessageListener and a JMS ExceptionListener, and to ensure that the IBM MQ classes for JMS are consistent with the JMS specification, the default value for the ASYNC\_EXCEPTIONS JMS ConnectionFactory property is changed to ASYNC\_EXCEPTIONS\_CONNECTIONBROKEN for the IBM MQ classes for JMS. As a result, by default, only exceptions corresponding to broken connection error codes are delivered to an application's JMS ExceptionListener.

**V9.0.0.1** APAR IT14820, included from IBM MQ Version 9.0.0, Fix Pack 1, updates IBM MQ classes for JMS so that:

- An ExceptionListener registered by an application is invoked for any connection broken exceptions, regardless of whether the application is using synchronous or asynchronous message consumers.
- An ExceptionListener registered by an application is invoked if a TCP/IP socket used by a JMS Session is broken.
- Non-connection broken exceptions (for example MQRC\_GET\_INHIBITED) that arise during message delivery are delivered to an application's ExceptionListener when the application is using asynchronous message consumers and the JMS ConnectionFactory used by the application has the ASYNC\_EXCEPTIONS property set to the value ASYNC\_EXCEPTIONS\_ALL.

**Note:** An ExceptionListener is only invoked once for a connection broken exception, even if two TCP/IP connections (one used by a JMS Connection and one used by a JMS Session) are broken.

For any other type of problem, a JMSException exception is thrown by the current JMS API call.

If an application does not register an exception listener with a Connection object, any exceptions that would have been delivered to the exception listener are written to the IBM MQ classes for JMS for JMS log.

### Related information:

Architecture and overview of features  
ASYNCEXCEPTION

### Accessing IBM MQ features from an IBM MQ classes for JMS application:

IBM MQ classes for JMS provides facilities to exploit a number of features of IBM MQ.

**Attention:** These features are outside the JMS specification or, in certain cases, violate the JMS specification. If you use them, your application is unlikely to be compatible with other JMS providers. Those features which do not comply with the JMS specification are labeled with an Attention notice.

*Reading and writing the message descriptor from an IBM MQ classes for JMS application:*

You control the ability to access the message descriptor (MQMD) by setting properties on a Destination and a Message.

Some IBM MQ applications require specific values to be set in the MQMD of messages sent to them. IBM MQ classes for JMS provides message attributes that allow JMS applications to set MQMD fields and so enable JMS applications to "drive" IBM MQ applications.

You must set the Destination object property `WMQ_MQMD_WRITE_ENABLED` to true for the setting of MQMD properties to have any effect. You can then use the property setting methods of the message (for example `setStringProperty`) to assign values to the MQMD fields. All MQMD fields are exposed except `StrucId` and `Version`; `BackoutCount` can be read but not written to.

This example results in a message being put to a queue or topic with `MQMD.UserIdentifier` set to "JoeBloggs".

```
// Create a ConnectionFactory, connection, session, producer, message
// ...

// Create a destination
// ...

// Enable MQMD write
dest.setBooleanProperty(WMQConstants.WMQ_MQMD_WRITE_ENABLED, true);

// Optionally, set a message context if applicable for this MD field
dest.setIntProperty(WMQConstants.WMQ_MQMD_MESSAGE_CONTEXT,
 WMQConstants.WMQ_MDCTX_SET_IDENTITY_CONTEXT);

// On the message, set property to provide custom UserId
msg.setStringProperty("JMS_IBM_MQMD_UserIdentifier", "JoeBloggs");

// Send the message
// ...
```

It is necessary to set `WMQ_MQMD_MESSAGE_CONTEXT` before setting `JMS_IBM_MQMD_UserIdentifier`. For more information about the use of `WMQ_MQMD_MESSAGE_CONTEXT`, see "JMS message object properties" on page 964.

Similarly, you can extract the contents of the MQMD fields by setting `WMQ_MQMD_READ_ENABLED` to true before receiving a message and then using the get methods of the message, such as `getStringProperty`. Any properties received are read-only.

This example results in the *value* field holding the value of the `MQMD.ApplIdentityData` field of a message got from a queue or a topic.

```
// Create a ConnectionFactory, connection, session, consumer
// ...

// Create a destination
// ...

// Enable MQMD read
dest.setBooleanProperty(WMQConstants.WMQ_MQMD_READ_ENABLED, true);

// Receive a message
// ...

// Get MQMD field value using a property
String value = rcvMsg.getStringProperty("JMS_IBM_MQMD_ApplIdentityData");
```

*JMS destination object properties:*

Two properties of the Destination object control access to the MQMD from JMS, and a third controls message context.

*Table 134. Property names and descriptions*

Property	Short form	Description
WMQ_MQMD_WRITE_ENABLED	MDW	Whether a JMS application can set the values of MQMD fields
WMQ_MQMD_READ_ENABLED	MDR	Whether a JMS application can extract the values of MQMD fields
WMQ_MQMD_MESSAGE_CONTEXT	MDCTX	What level of message context is to be set by the JMS application. The application must be running with appropriate context authority for this property to take effect

*Table 135. Property names, values, and set methods*

Property	Valid values in administration tool (defaults in bold)	Valid values in programs	Set method
WMQ_MQMD_WRITE_ENABLED	<ul style="list-style-type: none"> <li>• <b>NO</b> All JMS_IBM_MQMD* properties are ignored and their values are not copied into the underlying MQMD structure.</li> <li>• YES JMS_IBM_MQMD* properties are processed. Their values are copied into the underlying MQMD structure.</li> </ul>	<ul style="list-style-type: none"> <li>• <b>False</b></li> <li>• True</li> </ul>	setMQMDWriteEnabled
WMQ_MQMD_READ_ENABLED	<ul style="list-style-type: none"> <li>• <b>NO</b> When sending messages, the JMS_IBM_MQMD* properties on a sent message are not updated to reflect the updated field values in the MQMD. When receiving messages, none of the JMS_IBM_MQMD* properties are available on a received message, even if the sender had set some or all of them.</li> <li>• YES When sending messages, all of the JMS_IBM_MQMD* properties on a sent message are updated to reflect the updated field values in the MQMD, including those that the sender did not set explicitly. When receiving messages, all of the JMS_IBM_MQMD* properties are available on a received message, including those that the sender did not set explicitly.</li> </ul>	<ul style="list-style-type: none"> <li>• <b>False</b></li> <li>• True</li> </ul>	setMQMDReadEnabled

Table 135. Property names, values, and set methods (continued)

Property	Valid values in administration tool (defaults in bold)	Valid values in programs	Set method
WMQ_MQMD_MESSAGE_CONTEXT	<ul style="list-style-type: none"> <li>• <b>DEFAULT</b> The MQOPEN API call and the MQPMO structure specify no explicit message context options</li> <li>• SET_IDENTITY_CONTEXT The MQOPEN API call specifies the message context option MQOO_SET_IDENTITY_CONTEXT and the MQPMO structure specifies MQPMO_SET_IDENTITY_CONTEXT</li> <li>• SET_ALL_CONTEXT The MQOPEN API call specifies the message context option MQOO_SET_ALL_CONTEXT and the MQPMO structure specifies MQPMO_SET_ALL_CONTEXT</li> </ul>	<ul style="list-style-type: none"> <li>• <b>WMQ_MDCTX_DEFAULT</b></li> <li>• WMQ_MDCTX_SET_IDENTITY_CONTEXT</li> <li>• WMQ_MDCTX_SET_ALL_CONTEXT</li> </ul>	setMQMDMessageContext

JMS message object properties:

Message object properties prefixed JMS\_IBM\_MQMD allow you to set or read the corresponding MQMD field.

### Sending messages

All MQMD fields except StrucId and Version are represented. These properties refer only to the MQMD fields; where a property occurs both in the MQMD and in the MQRFH2 header, the version in the MQRFH2 is not set or extracted.

Any of these properties can be set, except JMS\_IBM\_MQMD\_BackoutCount. Any value set for JMS\_IBM\_MQMD\_BackoutCount is ignored.

If a property has a maximum length and you supply a value that is too long, the value is truncated.

For certain properties, you must also set the WMQ\_MQMD\_MESSAGE\_CONTEXT property on the Destination object. The application must be running with appropriate context authority for this property to take effect. If you do not set WMQ\_MQMD\_MESSAGE\_CONTEXT to an appropriate value, the property value is ignored. If you set WMQ\_MQMD\_MESSAGE\_CONTEXT to an appropriate value but you do not have sufficient context authority for the queue manager, a JMSEException is issued. Properties requiring specific values of WMQ\_MQMD\_MESSAGE\_CONTEXT are as follows.

The following properties require WMQ\_MQMD\_MESSAGE\_CONTEXT to be set to WMQ\_MDCTX\_SET\_IDENTITY\_CONTEXT or WMQ\_MDCTX\_SET\_ALL\_CONTEXT:

- JMS\_IBM\_MQMD\_UserIdentifier
- JMS\_IBM\_MQMD\_AccountingToken
- JMS\_IBM\_MQMD\_ApplIdentityData

The following properties require WMQ\_MQMD\_MESSAGE\_CONTEXT to be set to WMQ\_MDCTX\_SET\_ALL\_CONTEXT :

- JMS\_IBM\_MQMD\_PutApplType
- JMS\_IBM\_MQMD\_PutApplName

- JMS\_IBM\_MQMD\_PutDate
- JMS\_IBM\_MQMD\_PutTime
- JMS\_IBM\_MQMD\_ApplOriginData

## Receiving messages

All these properties are available on a received message if WMQ\_MQMD\_READ\_ENABLED property is set to true, irrespective of the actual properties the producing application has set. An application cannot modify the properties of a received message unless all properties are cleared first, according to the JMS specification. The received message can be forwarded without modifying the properties.

**Attention:** If your application receives a message from a destination with WMQ\_MQMD\_READ\_ENABLED property set to true, and forwards it to a destination with WMQ\_MQMD\_WRITE\_ENABLED set to true, this results in all the MQMD field values of the received message being copied into the forwarded message.

## Table of properties

This table lists the properties of the Message object representing the MQMD fields. See the links for full descriptions of the fields and their allowable values.

Table 136. Property names, descriptions, and types

Property	Description	Java Type	Link to full description
JMS_IBM_MQMD_Report	Options for report messages	Integer	Report
JMS_IBM_MQMD_MsgType	Message type	Integer	MsgType
JMS_IBM_MQMD_Expiry	Message lifetime	Integer	Expiry
JMS_IBM_MQMD_Feedback	Feedback or reason code	Integer	Feedback
JMS_IBM_MQMD_Encoding	Numeric encoding of message data	Integer	Encoding
JMS_IBM_MQMD_CodedCharSetId	Character set identifier of message data	Integer	CodedCharSetId
JMS_IBM_MQMD_Format	Format name of message data	String	Format
JMS_IBM_MQMD_Priority <sup>1</sup>	Message priority	Integer	Priority
JMS_IBM_MQMD_Persistence	Message persistence	Integer	Persistence
JMS_IBM_MQMD_MsgId <sup>2</sup>	Message identifier	Object (byte[]) <sup>4</sup>	MsgId
JMS_IBM_MQMD_CorrelId <sup>3</sup>	Correlation identifier	Object (byte[]) <sup>4</sup>	CorrelId
JMS_IBM_MQMD_BackoutCount	Backout counter	Integer	BackoutCount
JMS_IBM_MQMD_ReplyToQ	Name of reply queue	String	ReplyToQ
JMS_IBM_MQMD_ReplyToQMgr	Name of reply queue manager	String	ReplyToQMgr
JMS_IBM_MQMD_UserIdentifier	User identifier	String	UserIdentifier
JMS_IBM_MQMD_AccountingToken	Accounting token	Object (byte[]) <sup>4</sup>	AccountingToken
JMS_IBM_MQMD_ApplIdentityData	Application data relating to identity	String	ApplIdentityData
JMS_IBM_MQMD_PutApplType	Type of application that put the message	Integer	PutApplType

Table 136. Property names, descriptions, and types (continued)

Property	Description	Java Type	Link to full description
JMS_IBM_MQMD_PutApplName	Name of application that put the message	String	PutApplName
JMS_IBM_MQMD_PutDate	Date when message was put	String	PutDate
JMS_IBM_MQMD_PutTime	Time when message was put	String	PutTime
JMS_IBM_MQMD_ApplOriginData	Application data relating to origin	String	ApplOriginData
JMS_IBM_MQMD_GroupId	Group identifier	Object (byte[]) <sup>4</sup>	GroupId
JMS_IBM_MQMD_MsgSeqNumber	Sequence number of logical message within group	Integer	MsgSeqNumber
JMS_IBM_MQMD_Offset	Offset of data in physical message from start of logical message	Integer	Offset
JMS_IBM_MQMD_MsgFlags	Message flags	Integer	MsgFlags
JMS_IBM_MQMD_OriginalLength	Length of original message	Integer	OriginalLength

1. **Attention:** If you assign a value to JMS\_IBM\_MQMD\_Priority that is not within the range 0-9, this violates the JMS specification.
2. **Attention:** The JMS specification states that the message ID must be set by the JMS provider and that it must either be unique or null. If you assign a value to JMS\_IBM\_MQMD\_MsgId, this value is copied to the JMSMessageID. Thus it is not set by the JMS provider and might not be unique: this violates the JMS specification.
3. **Attention:** If you assign a value to JMS\_IBM\_MQMD\_CorrelId that starts with the string 'ID:', this violates the JMS specification.
4. **Attention:** The use of byte array properties on a message violates the JMS specification.

*Accessing IBM MQ Message data from an application using IBM MQ classes for JMS:*

You can access the complete IBM MQ message data within an application using IBM MQ classes for JMS. To access all the data, the message must be a JMSBytesMessage. The body of the JMSBytesMessage includes any MQRFH2 header, any other IBM MQ headers, and the following message data.

Set the WMQ\_MESSAGE\_BODY property of the destination to WMQ\_MESSAGE\_BODY\_MQ, to receive all the message body data in the JMSBytesMessage.

If WMQ\_MESSAGE\_BODY is set to WMQ\_MESSAGE\_BODY\_JMS or WMQ\_MESSAGE\_BODY\_UNSPECIFIED, the message body is returned without the JMS MQRFH2 header, and the properties of the JMSBytesMessage reflect the properties set in the RFH2.

Some applications cannot use the functions described in this topic. If an application is connected to an IBM MQ V6 queue manager, or if it has set PROVIDERVERSION to 6, the functions are not available.

### **Sending a message**

When sending messages the destination property, WMQ\_MESSAGE\_BODY, takes precedence over WMQ\_TARGET\_CLIENT.

If `WMQ_MESSAGE_BODY` is set to `WMQ_MESSAGE_BODY_JMS`, IBM MQ classes for JMS automatically generates an MQRFH2 header based on the settings of the `JMSMessage` properties and header fields.

If `WMQ_MESSAGE_BODY` is set to `WMQ_MESSAGE_BODY_MQ`, no additional header is added to the message body

If `WMQ_MESSAGE_BODY` is set to `WMQ_MESSAGE_BODY_UNSPECIFIED`, IBM MQ classes for JMS sends an MQRFH2 header, unless `WMQ_TARGET_CLIENT` is set to `WMQ_TARGET_DEST_MQ`. On receive, setting `WMQ_TARGET_CLIENT` to `WMQ_TARGET_DEST_MQ` results in any MQRFH2 being removed from the message body.

**Note:** `JMSBytesMessage` and `JMSTextMessage` do not require an MQRFH2, whereas `JMSStreamMessage`, `JMSMapMessage`, and `JMSObjectMessage` do.

`WMQ_MESSAGE_BODY_UNSPECIFIED` is the default setting for `WMQ_MESSAGE_BODY`, and `WMQ_TARGET_DEST_JMS` is the default setting for `WMQ_TARGET_CLIENT`.

If you send a `JMSBytesMessage`, you can override the default settings for the JMS message body when the IBM MQ message is constructed. Use the following properties:

- `JMS_IBM_Format` or `JMS_IBM_MQMD_Format`: This property specifies the format of the IBM MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.
- `JMS_IBM_Character_Set` or `JMS_IBM_MQMD_CodedCharSetId`: This property specifies the CCSID of the IBM MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.
- `JMS_IBM_Encoding` or `JMS_IBM_MQMD_Encoding`: This property specifies the encoding of the IBM MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.

If both types of property are specified, the `JMS_IBM_MQMD_*` properties override the corresponding `JMS_IBM_*` properties, as long as the destination property `WMQ_MQMD_WRITE_ENABLED` is set to true.

The differences in effect between setting message properties using `JMS_IBM_MQMD_*` and `JMS_IBM_*` are significant:

1. The `JMS_IBM_MQMD_*` properties are specific to the IBM MQ JMS provider.
2. The `JMS_IBM_MQMD_*` properties are only set in the MQMD. `JMS_IBM_*` properties are set in the MQMD only if the message does not have an MQRFH2 JMS header. Otherwise they are set in the JMS RFH2 header.
3. The `JMS_IBM_MQMD_*` properties have no affect on the encoding of text and numbers written into a `JMSMessage`.

A receiving application is likely to assume the values of `MQMD.Encoding` and `MQMD.CodedCharSetId` correspond to the encoding and character set of numbers and text in the message body. If `JMS_IBM_MQMD_*` properties are used, it is the responsibility of the sending application to make it so. The encoding and character set of numbers and text in the message body are set by the `JMS_IBM_*` properties.

The badly coded snippet in Figure 143 on page 968 sends a message encoded in character set 1208, with `MQMD.CodedCharSetId` set to 37.

---

a. Send wrongly encoded message

```
TextMessage tmo = session.createTextMessage();
((MQDestination) destination).setMessageBodyStyle
 (WMQConstants.WMQ_MESSAGE_BODY_MQ);
((MQDestination) destination).setMQMDWriteEnabled(true);
tmo.setIntProperty(WMQConstants.JMS_IBM_MQMD_CODEDCHARSETID, 37);
tmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 1208);
tmo.setText("String one");
producer.send(tmo);
```

b. Receiving the message, relying on the value of JMS\_IBM\_CHARACTER\_SET set by the value of MQMD.CodedCharSetId:

```
TextMessage tmi = (TextMessage) cons.receive();
System.out.println("Message is \" + tmi.getText() + "\"");
```

c. Resulting output:

```
Message is "éËË!>...??>?"
```

---

*Figure 143. Inconsistently coded MQMD and message data*

Either of the snippets of code in Figure 144 results in a message being put to a queue or topic, with its body containing the application payload without an automatically generated MQRFH2 header being added.

---

1. Setting WMQ\_MESSAGE\_BODY\_MQ:

```
((MQDestination) destination).setMessageBodyStyle
 (WMQConstants.WMQ_MESSAGE_BODY_MQ);
```

2. Setting WMQ\_TARGET\_DEST\_MQ:

```
((MQDestination) destination).setMessageBodyStyle
 (WMQConstants.WMQ_MESSAGE_BODY_UNSPECIFIED);
((MQDestination) destination).
 setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ);
```

---

*Figure 144. Send a message with an MQ message body.*

## Receiving a message

If WMQ\_MESSAGE\_BODY is set to WMQ\_MESSAGE\_BODY\_JMS, the inbound JMS message type and body are determined by the contents of the received WebSphere MQ message. The message type and body are determined by fields in the MQRFH2 header, or in the MQMD, if there is no MQRFH2.

If WMQ\_MESSAGE\_BODY is set to WMQ\_MESSAGE\_BODY\_MQ, the inbound JMS message type is JMSBytesMessage. The JMS message body is the message data returned by the underlying MQGET API call. The length of message body is the length returned by the MQGET call. The character set and encoding of the data in the message body is determined by the CodedCharSetId and Encoding fields of the MQMD. The format of the data in the message body is determined by the Format field of the MQMD

If WMQ\_MESSAGE\_BODY is set to WMQ\_MESSAGE\_BODY\_UNSPECIFIED, the default value, IBM MQ classes for JMS sets it to WMQ\_MESSAGE\_BODY\_JMS.

When you receive a JMSBytesMessage, you can decode it by reference to the following properties:

- JMS\_IBM\_Format or JMS\_IBM\_MQMD\_Format: This property specifies the format of the IBM MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.
- JMS\_IBM\_Character\_Set or JMS\_IBM\_MQMD\_CodedCharSetId: This property specifies the CCSID of the IBM MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.



- `JMS_IBM_Encoding` or `JMS_IBM_MQMD_Encoding`: This property specifies the encoding of the IBM MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.

The following code snippet results in a received message that is a `JMSBytesMessage`. Irrespective of the content of the received message and of the format field of the received MQMD, the message is a `JMSBytesMessage`.

```
((MQDestination)destination).setMessageBodyStyle
(WMQConstants.WMQ_MESSAGE_BODY_MQ);
```

*Destination property* `WMQ_MESSAGE_BODY`:

`WMQ_MESSAGE_BODY` determines whether a JMS application processes the MQRFH2 of an IBM MQ message as part of the message payload (that is, as part of the JMS message body).

*Table 137. Property names and descriptions*

Property	Short form	Description
<code>WMQ_MESSAGE_BODY</code>	<code>MBODY</code>	Whether a JMS application processes the MQRFH2 of an IBM MQ message as part of the message payload (that is, as part of the JMS message body).

*Table 138. Property names, values, and set methods*

Property	Valid values in administration tool (defaults in bold)	Valid values in programs	Set method
<code>WMQ_MESSAGE_BODY</code>	<ul style="list-style-type: none"> <li>• <b>UNSPECIFIED</b> When sending, IBM MQ classes for JMS does or does not generate and include an MQRFH2 header, depending on the value of <code>WMQ_TARGET_CLIENT</code>. When receiving, acts as value JMS.</li> <li>• JMS When sending, IBM MQ classes for JMS automatically generates an MQRFH2 header and includes it in the IBM MQ message. When receiving, IBM MQ classes for JMS set the JMS message properties according to values in the MQRFH2 (if present); it does not present the MQRFH2 as part of the JMS message body.</li> <li>• MQ When sending, IBM MQ classes for JMS does not generate an MQRFH2. When receiving, IBM MQ classes for JMS presents the MQRFH2 as part of the JMS message body.</li> </ul>	<ul style="list-style-type: none"> <li>• <b>WMQ_MESSAGE_BODY_UNSPECIFIED</b></li> <li>• <code>WMQ_MESSAGE_BODY_JMS</code></li> <li>• <code>WMQ_MESSAGE_BODY_MQ</code></li> </ul>	<code>setMessageBodyStyle</code>

*JMS persistent messages:*

IBM MQ classes for JMS applications can use the **NonPersistentMessageClass** queue attribute to provide better performance for JMS persistent messages, at the expense of some reliability.

An IBM MQ queue has an attribute called **NonPersistentMessageClass**. The value of this attribute determines whether nonpersistent messages on the queue are discarded when the queue manager restarts.

You can set the attribute for a local queue by using the IBM MQ Script (MQSC) command, **DEFINE QLOCAL**, with either of the following parameters:

**NPMCLASS(NORMAL)**

Nonpersistent messages on the queue are discarded when the queue manager restarts. This is the default value.

**NPMCLASS(HIGH)**

Nonpersistent messages on the queue are not discarded when the queue manager restarts following a quiesced or immediate shutdown. Nonpersistent messages might be discarded, however, following a preemptive shutdown or a failure.

This topic describes how IBM MQ classes for JMS applications can use this queue attribute to provide better performance for JMS persistent messages.

The **PERSISTENCE** property of a Queue or Topic object can have the value **HIGH**. You can use the IBM MQ JMS administration tool to set this value, or an application can call the `Destination.setPersistence()` method passing the value `WMQConstants.WMQ_PER_NPHIGH` as a parameter.

If an application sends a JMS persistent message or a JMS nonpersistent message to a destination where the **PERSISTENCE** property has the value **HIGH**, and the underlying IBM MQ queue is set to **NPMCLASS(HIGH)**, the message is put on the queue as an IBM MQ nonpersistent message. If the **PERSISTENCE** property of the destination does not have the value **HIGH**, or if the underlying queue is set to **NPMCLASS(NORMAL)**, a JMS persistent message is put on the queue as an IBM MQ persistent message, and a JMS nonpersistent message is put on the queue as an IBM MQ nonpersistent message.

If a JMS persistent message is put on a queue as an IBM MQ nonpersistent message, and you want to ensure that the message is not discarded following a quiesced or immediate shutdown of a queue manager, all queues through which the message might be routed must be set to **NPMCLASS(HIGH)**. In the publish/subscribe domain, these queues include subscriber queues. As an aid to enforcing this configuration, IBM MQ classes for JMS throws an `InvalidDestinationException` if an application tries to create a message consumer for a destination where the **PERSISTENCE** property has the value **HIGH** and the underlying IBM MQ queue is set to **NPMCLASS(NORMAL)**.

Setting the **PERSISTENCE** property of a destination to **HIGH** does not affect how a message is received from that destination. A message sent as a JMS persistent message is received as a JMS persistent message, and a message sent as a JMS nonpersistent message is received as a JMS nonpersistent message.

When an application sends the first message to a destination where the **PERSISTENCE** property has the value **HIGH**, or when an application creates the first message consumer for a destination where the **PERSISTENCE** property has the value **HIGH**, IBM MQ classes for JMS issues an `MQINQ` call to determine whether **NPMCLASS(HIGH)** is set on the underlying IBM MQ queue. The application must therefore have the authority to inquire on the queue. In addition, IBM MQ classes for JMS preserves the result of the `MQINQ` call until the destination is deleted, and does not issue more `MQINQ` calls. Therefore, if you change the **NPMCLASS** setting on the underlying queue while the application is still using the destination, IBM MQ classes for JMS does not notice the new setting.

By allowing JMS persistent messages to be put on IBM MQ queues as IBM MQ nonpersistent messages, you are gaining performance at the expense of some reliability. If you require maximum reliability for JMS persistent messages, do not send the messages to a destination where the PERSISTENCE property has the value HIGH.

The JMS Layer can use SYSTEM.JMS.TEMPQ.MODEL, instead of SYSTEM.DEFAULT.MODEL.QUEUE. SYSTEM.JMS.TEMPQ.MODEL creates permanent dynamic queues that accept persistent messages, because SYSTEM.DEFAULT.MODEL.QUEUE cannot accept persistent messages. To use temporary queues to accept persistent messages, you must therefore use SYSTEM.JMS.TEMPQ.MODEL, or change the model queue to an alternative queue of your choosing.

*Using TLS with IBM MQ classes for JMS:*

IBM MQ classes for JMS applications can use Transport Layer Security (TLS) encryption. To do this they require a JSSE provider.

IBM MQ classes for JMS connections using TRANSPORT(CLIENT) support TLS encryption. TLS provides communication encryption, authentication, and message integrity. It is typically used to secure communications between any two peers on the Internet or within an intranet.

IBM MQ classes for JMS uses Java Secure Socket Extension (JSSE) to handle TLS encryption, and therefore requires a JSSE provider. JSE v1.4 JVMs have a JSSE provider built-in. Details of how to manage and store certificates can vary from provider to provider. For information about this, see your JSSE provider's documentation.

This section assumes that your JSSE provider is correctly installed and configured, and that suitable certificates have been installed and made available to your JSSE provider. You can now use JMSAdmin to set a number of administrative properties.

If your IBM MQ classes for JMS application uses a client channel definition table (CCDT) to connect to a queue manager, see “Using a client channel definition table with IBM MQ classes for JMS” on page 982.

*SSLCIPHERSUITE object property:*

Set SSLCIPHERSUITE to enable TLS encryption on a ConnectionFactory object.

To enable TLS encryption on a ConnectionFactory object, use JMSAdmin to set the SSLCIPHERSUITE property to a CipherSuite supported by your JSSE provider. This must match the CipherSpec set on the target channel. However, CipherSuites are distinct from CipherSpecs and therefore have different names. “TLS CipherSpecs and CipherSuites in IBM MQ classes for JMS” on page 975 contains a table mapping the CipherSpecs supported by IBM MQ to their equivalent CipherSuites as known to JSSE. For more information about CipherSpecs and CipherSuites with IBM MQ, see Securing.

For example, to set up a ConnectionFactory object that can be used to create a connection over an TLS enabled MQI channel with a CipherSpec of TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA, issue the following command to JMSAdmin:

```
ALTER CF(my.cf) SSLCIPHERSUITE(SSL_RSA_WITH_AES_128_CBC_SHA)
```

This can also be set from an application, using the setSSLCipherSuite() method on an MQConnectionFactory object.

For convenience, if a CipherSpec is specified on the SSLCIPHERSUITE property, JMSAdmin attempts to map the CipherSpec to an appropriate CipherSuite and issues a warning. This attempt to map is not made if the property is specified by an application.

Alternatively, use the Client Channel Definition Table (CCDT). For more information, see “Using a client channel definition table with IBM MQ classes for JMS” on page 982.

*SSLFIPSREQUIRED object property:*

If you require a connection to use a CipherSuite that is supported by the IBM Java JSSE FIPS provider (IBMJSSEFIPS), set the SSLFIPSREQUIRED property of the connection factory to YES.

The default value of this property is NO, which means that a connection can use any CipherSuite that is supported by IBM MQ.

If an application uses more than one connection, the value of SSLFIPSREQUIRED that is used when the application creates the first connection determines the value that is used when the application creates any subsequent connection. This means that the value of the SSLFIPSREQUIRED property of the connection factory that is used to create a subsequent connection is ignored. You must restart the application if you want to use a different value of SSLFIPSREQUIRED.

An application can set this property by calling the setSSLFipsRequired() method of a ConnectionFactory object. The property is ignored if no CipherSuite is set.

**Related information:**

Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client  
Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows

*SSLPEERNAME object property:*

Use SSLPEERNAME to specify a distinguished name pattern, to ensure that your JMS application connects to the correct queue manager.

A JMS application can ensure that it connects to the correct queue manager by specifying a distinguished name (DN) pattern. The connection succeeds only if the queue manager presents a DN that matches the pattern. For more details of the format of this pattern, see the related topics.

The DN is set using the SSLPEERNAME property of a ConnectionFactory object. For example, the following JMSAdmin command sets a ConnectionFactory object to expect the queue manager to identify itself with a Common Name beginning with the characters QMGR., and with at least two Organizational Unit names, the first of which must be IBM and the second WEBSHERE:

```
ALTER CF(my.cf) SSLPEERNAME(CN=QMGR.*, OU=IBM, OU=WEBSHERE)
```

Checking is not case sensitive and semicolons can be used in place of commas. SSLPEERNAME can also be set from an application using the setSSLPeerName() method on an MQConnectionFactory object. If this property is not set, no checking is performed on the Distinguished Name supplied by the queue manager. This property is ignored if no CipherSuite is set.

*SSLCERTSTORES* object property:

Use SSLCERTSTORES to specify a list of LDAP servers to use for certificate revocation list (CRL) checking.

It is common to use a certificate revocation list (CRL) to identify certificates that are no longer trusted. CRLs are typically hosted on LDAP servers. JMS allows an LDAP server to be specified for CRL checking under Java 2 v1.4 or later. The following JMSAdmin example directs JMS to use a CRL hosted on an LDAP server named crl1.ibm.com:

```
ALTER CF(my.cf) SSLCRL(ldap://crl1.ibm.com)
```

**Note:** To use a CertStore successfully with a CRL hosted on an LDAP server, make sure that your Java Software Development Kit (SDK) is compatible with the CRL. Some SDKs require that the CRL conforms to RFC 2587, which defines a schema for LDAP v2. Most LDAP v3 servers use RFC 2256 instead.

If your LDAP server is not running on the default port of 389, you can specify the port by appending a colon (:) and the port number to the host name. If the certificate presented by the queue manager is present in the CRL hosted on crl1.ibm.com, the connection is not completed. To avoid a single point of failure, JMS allows multiple LDAP servers to be supplied by supplying a list of LDAP servers delimited by the space character. Here is an example:

```
ALTER CF(my.cf) SSLCRL(ldap://crl1.ibm.com ldap://crl2.ibm.com)
```

When multiple LDAP servers are specified, JMS tries each one in turn until it finds a server with which it can successfully verify the queue manager's certificate. Each server must contain identical information.

A string in this format can be supplied by an application on the MQConnectionFactory.setSSLCertStores() method. Alternatively, the application can create one or more java.security.cert.CertStore objects, place these in a suitable Collection object, and supply this Collection object to the setSSLCertStores() method. In this way, the application can customize CRL checking. See your JSSE documentation for details on constructing and using CertStore objects.

The certificate presented by the queue manager when a connection is being set up is validated as follows:

1. The first CertStore object in the Collection identified by sslCertStores is used to identify a CRL server.
2. An attempt is made to contact the CRL server.
3. If the attempt is successful, the server is searched for a match for the certificate.
  - a. If the certificate is found to be revoked, the search process is over and the connection request fails with reason code MQRC\_SSL\_CERTIFICATE\_REVOKED.
  - b. If the certificate is not found, the search process is over and the connection is allowed to proceed.
4. If the attempt to contact the server is unsuccessful, the next CertStore object is used to identify a CRL server and the process repeats from step 2.

If this was the last CertStore in the Collection, or if the Collection contains no CertStore objects, the search process has failed and the connection request fails with reason code MQRC\_SSL\_CERT\_STORE\_ERROR.

The Collection object determines the order in which CertStores are used.

If your application uses setSSLCertStores() to set a Collection of CertStore objects, the MQConnectionFactory can no longer be bound into a JNDI namespace. Attempting to do so causes an exception. If the sslCertStores property is not set, no revocation checking is performed on the certificate provided by the queue manager. This property is ignored if no CipherSuite is set.

*SSLRESETCOUNT object property:*

This property represents the total number of bytes sent and received by a connection before the secret key that is used for encryption is renegotiated.

The number of bytes sent is the number before encryption, and the number of bytes received is the number after decryption. The number of bytes also includes control information sent and received by IBM MQ classes for JMS.

For example, to configure a `ConnectionFactory` object that can be used to create a connection over an TLS enabled MQI channel with a secret key that is renegotiated after 4 MB of data have flowed, issue the following command to JMSAdmin:

```
ALTER CF(my.cf) SSLRESETCOUNT(4194304)
```

An application can set this property by calling the `setSSLResetCount()` method of a `ConnectionFactory` object.

If the value of this property is zero, which is the default value, the secret key is never renegotiated. The property is ignored if no `CipherSuite` is set.

*SSLSocketFactory object property:*

To customize other aspects of the TLS connection for an application, create an `SSLSocketFactory` and configure JMS to use it.

You might want to customize other aspects of the TLS connection for an application. For example, you might want to initialize cryptographic hardware or change the keystore and truststore in use. To do this, the application must first create a `javax.net.ssl.SSLSocketFactory` object that is customized accordingly. See your JSSE documentation for information about how to do this, because the customizable features vary from provider to provider. After a suitable `SSLSocketFactory` object is obtained, use the `MQConnectionFactory.setSSLSocketFactory()` method to configure JMS to use the customized `SSLSocketFactory` object.

If your application uses the `setSSLSocketFactory()` method to set a customized `SSLSocketFactory` object, the `MQConnectionFactory` object can no longer be bound into a JNDI namespace. Attempting to do so causes an exception. If this property is not set, the default `SSLSocketFactory` object is used. See your JSSE documentation for details of the behavior of the default `SSLSocketFactory` object. This property is ignored if no `CipherSuite` is set.

**Important:** Do not assume that the use of the SSL properties ensures security when a `ConnectionFactory` object is retrieved from a JNDI namespace that is not itself secure. Specifically, the standard LDAP implementation of JNDI is not secure. An attacker can imitate the LDAP server, misleading a JMS application into connecting to the wrong server without noticing. With suitable security arrangements in place, other implementations of JNDI (such as the `fscontext` implementation) are secure.

*Making changes to the JSSE keystore or truststore:*

If you make changes to the keystore or truststore, you must take certain actions for the changes to be picked up.

If you change the contents of the JSSE keystore or truststore, or change the location of the keystore or truststore file, IBM MQ classes for JMS applications that are running at the time do not automatically pick up the changes. For the changes to take effect, the following actions must be performed:

- The applications must close all their connections, and destroy any unused connections in connection pools.
- If your JSSE provider caches information from the keystore and truststore, this information must be refreshed.

After these actions have been performed, the applications can then re-create their connections.

Depending on how you design your applications, and on the function provided by your JSSE provider, it might be possible to perform these actions without stopping and restarting your applications. However, stopping and restarting the applications might be the simplest solution.

*TLS CipherSpecs and CipherSuites in IBM MQ classes for JMS:*

The ability of IBM MQ classes for JMS applications to establish connections to a queue manager, depends on the CipherSpec specified at the server end of the MQI channel and the CipherSuite specified at the client end.

The following table lists the CipherSpecs supported by IBM MQ and their equivalent CipherSuites.

You should review the topic *Deprecated CipherSpecs* to see if any of the CipherSpecs, listed in the following table, have been deprecated by IBM MQ and, if so, at which update the CipherSpec was deprecated.

**Important:** The CipherSuites listed are those supported by the IBM Java Runtime Environment (JRE) supplied with IBM MQ. The CipherSuites that are listed include those supported by the Oracle Java JRE. For more information about configuring your application to use an Oracle Java JRE, see *Configuring your application to use IBM Java or Oracle Java CipherSuite mappings*.

The table also indicates the protocol that is used for the communication, and whether or not the CipherSuite conforms to the FIPS 140-2 standard.

Ciphersuites denoted as FIPS 140-2 compliant can be used if the application has not been configured to enforce FIPS 140-2 compliance, but if FIPS 140-2 compliance has been configured for the application (see the following notes on configuration) only those CipherSuites which are marked as FIPS 140-2 compatible can be configured; attempting to use other CipherSuites results in an error.

**Note:** Each JRE can have multiple cryptographic security providers, each of which can contribute an implementation of the same CipherSuite. However, not all security providers are FIPS 140-2 certified. If FIPS 140-2 compliance is not enforced for an application then it is possible that an uncertified implementation of the CipherSuite might be used. Uncertified implementations might not operate in compliance with FIPS 140-2, even if the CipherSuite theoretically meets the minimum security level required by the standard. See the following notes for more information about configuring FIPS 140-2 enforcement in IBM MQ JMS applications.

For more information about FIPS 140-2 and Suite-B compliance for CipherSpecs and CipherSuites, see *Specifying CipherSpecs*. You might also need to be aware of information that concerns US Federal Information Processing Standards.

To use the full set of CipherSuites and to operate with certified FIPS 140-2 and/or Suite-B compliance, a suitable JRE is required. IBM Java 7 Service Refresh 4 Fix Pack 2 or a higher level of IBM JRE provides the appropriate support.

**Note:** To use some CipherSuites, the 'unrestricted' policy files need to be configured in the JRE. For more details of how policy files are set up in an SDK or JRE, see the *IBM SDK Policy files* topic in the *Security Reference for IBM SDK, Java Technology Edition* for the version you are using..

Table 139. CipherSpecs supported by IBM MQ and their equivalent CipherSuites

CipherSpec	Equivalent CipherSuite (IBM JRE)	Equivalent CipherSuite (Oracle JRE)	Protocol	FIPS 140-2 compatible
ECDHE_ECDSA_3DES_EDE_CBC_SHA256	SSL_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA256	TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA256	TLS	Yes
ECDHE_ECDSA_AES_128_CBC_SHA256	SSL_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLS	Yes
ECDHE_ECDSA_AES_128_GCM_SHA256	SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLS	Yes
ECDHE_ECDSA_AES_256_CBC_SHA384	SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLS	Yes
ECDHE_ECDSA_AES_256_GCM_SHA384	SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLS	Yes
ECDHE_ECDSA_NULL_SHA256	SSL_ECDHE_ECDSA_WITH_NULL_SHA	TLS_ECDHE_ECDSA_WITH_NULL_SHA	TLS	Yes
ECDHE_ECDSA_RC4_128_SHA256	SSL_ECDHE_ECDSA_WITH_RC4_128_SHA	TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	TLS	Yes
ECDHE_RSA_3DES_EDE_CBC_SHA256	SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA256	TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA256	TLS	Yes
ECDHE_RSA_AES_128_CBC_SHA256	SSL_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLS	Yes
ECDHE_RSA_AES_128_GCM_SHA256	SSL_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLS	Yes
ECDHE_RSA_AES_256_CBC_SHA384	SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLS	Yes
ECDHE_RSA_AES_256_GCM_SHA384	SSL_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLS	Yes
ECDHE_RSA_NULL_SHA256	SSL_ECDHE_RSA_WITH_NULL_SHA	TLS_ECDHE_RSA_WITH_NULL_SHA	TLS	Yes
ECDHE_RSA_RC4_128_SHA256	SSL_ECDHE_RSA_WITH_RC4_128_SHA	TLS_ECDHE_RSA_WITH_RC4_128_SHA	TLS	Yes
TLS_RSA_WITH_3DES_EDE_CBC_SHA <sub>1</sub>	SSL_RSA_WITH_3DES_EDE_CBC_SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	TLS	Yes
TLS_RSA_WITH_AES_128_CBC_SHA	SSL_RSA_WITH_AES_128_CBC_SHA	TLS_RSA_WITH_AES_128_CBC_SHA	TLS	Yes
TLS_RSA_WITH_AES_128_CBC_SHA256	SSL_RSA_WITH_AES_128_CBC_SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256	TLS	Yes
TLS_RSA_WITH_AES_128_GCM_SHA256	SSL_RSA_WITH_AES_128_GCM_SHA256	TLS_RSA_WITH_AES_128_GCM_SHA256	TLS	Yes
TLS_RSA_WITH_AES_256_CBC_SHA	SSL_RSA_WITH_AES_256_CBC_SHA	TLS_RSA_WITH_AES_256_CBC_SHA	TLS	Yes
TLS_RSA_WITH_AES_256_CBC_SHA256	SSL_RSA_WITH_AES_256_CBC_SHA256	TLS_RSA_WITH_AES_256_CBC_SHA256	TLS	Yes
TLS_RSA_WITH_AES_256_GCM_SHA384	SSL_RSA_WITH_AES_256_GCM_SHA384	TLS_RSA_WITH_AES_256_GCM_SHA384	TLS	Yes
TLS_RSA_WITH_DES_CBC_SHA	SSL_RSA_WITH_DES_CBC_SHA	SSL_RSA_WITH_DES_CBC_SHA	SSL	No
TLS_RSA_WITH_NULL_SHA256	SSL_RSA_WITH_NULL_SHA256	TLS_RSA_WITH_NULL_SHA256	TLS	Yes
TLS_RSA_WITH_RC4_128_SHA256	SSL_RSA_WITH_RC4_128_SHA	SSL_RSA_WITH_RC4_128_SHA	SSL	No

**Notes:**

1. This CipherSpec TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA is deprecated. However, it can still be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, you need to either avoid using triple DES, or enable secret key reset when using this CipherSpec.



## Configuring Ciphersuites and FIPS-compliance in an IBM MQ classes for JMS application

- An application that uses IBM MQ classes for JMS can use either of two methods to set the CipherSuite for a connection:
  - Call the setSSLCipherSuite method of a ConnectionFactory object.
  - Use the IBM MQ JMS administration tool to set the SSLCIPHERSUITE property of a ConnectionFactory object.
- An application that uses IBM MQ classes for JMS can use either of two methods to enforce FIPS 140-2 compliance:
  - Call the setSSLFipsRequired method of a ConnectionFactory object.
  - Use the IBM MQ JMS administration tool to set the SSLFIPSREQUIRED property of a ConnectionFactory object.

## Configuring your application to use IBM Java or Oracle Java CipherSuite mappings

You can configure whether your application uses the default IBM Java CipherSuite to IBM MQ CipherSpec mappings, or the Oracle CipherSuite to IBM MQ CipherSpec mappings. Therefore, you can use TLS CipherSuites whether your application uses an IBM JRE or an Oracle JRE. The Java System Property `com.ibm.mq.cfg.useIBMCipherMappings` controls which mappings are used. The property can be one of the following values:

**true** Use the IBM Java CipherSuite to IBM MQ CipherSpec mappings.  
This value is the default value.

**false** Use the Oracle CipherSuite to IBM MQ CipherSpec mappings.

For more information about using IBM MQ Java and TLS Ciphers, see the MQdev blog posts MQ Java, TLS Ciphers, Non-IBM JREs & APARs IT06775, IV66840, IT09423, IT10837, and The relationship between MQ CipherSpecs and Java Cipher Suites.

## Interoperability limitations

Certain CipherSuites might be compatible with more than one IBM MQ CipherSpec, depending on the protocol in use. However, only the CipherSuite/CipherSpec combination that uses the TLS version specified in Table 1 is supported. Attempting to use the unsupported combinations of CipherSuites and CipherSpecs will fail with an appropriate exception. Installations using any of these CipherSuite/CipherSpec combinations should move to a supported combination.

The following table shows the CipherSuites to which this limitation applies.

Table 140. CipherSuites and their supported and unsupported CipherSpecs

CipherSuite	Supported TLS CipherSpec	Unsupported SSL CipherSpec
SSL_RSA_WITH_3DES_EDE_CBC_SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA <sup>1</sup>	TRIPLE_DES_SHA_US
SSL_RSA_WITH_DES_CBC_SHA	TLS_RSA_WITH_DES_CBC_SHA	DES_SHA_EXPORT
SSL_RSA_WITH_RC4_128_SHA	TLS_RSA_WITH_RC4_128_SHA256	RC4_SHA_US

### Note:

1. This CipherSpec `TLS_RSA_WITH_3DES_EDE_CBC_SHA` is deprecated. However, it can still be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, you need to either avoid using triple DES, or enable secret key reset when using this CipherSpec.

*Writing channel exits in Java for IBM MQ classes for JMS:*

You create channel exits by defining Java classes that implement specified interfaces.

Three interfaces are defined in the `com.ibm.mq.exits` package:

- `WMQSendExit`, for a send exit
- `WMQReceiveExit`, for a receive exit
- `WMQSecurityExit`, for a security exit

The following sample code defines a class that implements all three interfaces:

```
public class MyMQExits implements
WMQSendExit, WMQReceiveExit, WMQSecurityExit {
 // Default constructor
 public MyMQExits(){
 }
 // This method implements the send exit interface
 public ByteBuffer channelSendExit(
 MQCXP channelExitParms,
 MQCD channelDefinition,
 ByteBuffer agentBuffer)
 {
 // Complete the body of the send exit here
 }
 // This method implements the receive exit interface
 public ByteBuffer channelReceiveExit(
 MQCXP channelExitParms,
 MQCD channelDefinition,
 ByteBuffer agentBuffer)
 {
 // Complete the body of the receive exit here
 }
 // This method implements the security exit interface
 public ByteBuffer channelSecurityExit(
 MQCXP channelExitParms,
 MQCD channelDefinition,
 ByteBuffer agentBuffer)
 {
 // Complete the body of the security exit here
 }
}
```

Each exit receives as parameters an MQCXP object and an MQCD object. These objects represent the MQCXP and MQCD structures defined in the procedural interface.

When a send exit is called, the `agentBuffer` parameter contains the data that is about to be sent to the server queue manager. A length parameter is not required because the expression `agentBuffer.limit()` provides the length of the data. The send exit returns as its value the data to be sent to the server queue manager. However, if the send exit is not the last send exit in a sequence of send exits, the data returned is passed instead to the next send exit in the sequence. A send exit can return a modified version of the data that it receives in the `agentBuffer` parameter, or it can return the data unchanged. The simplest possible exit body is therefore:

```
{ return agentBuffer; }
```

When a receive exit is called, the `agentBuffer` parameter contains the data that has been received from the server queue manager. The receive exit returns as its value the data to be passed to the application by IBM MQ classes for JMS. However, if the receive exit is not the last receive exit in a sequence of receive exits, the data returned is passed instead to the next receive exit in the sequence.

When a security exit is called, the `agentBuffer` parameter contains the data that has been received in a security flow from the security exit at the server end of the connection. The security exit returns as its value the data to be sent in a security flow to the server security exit.

Channel exits are called with a buffer that has a backing array. For best performance, the exit should return a buffer with a backing array.

Up to 32 characters of user data can be passed to a channel exit when it is called. The exit accesses the user data by calling the `getExitData()` method of the MQCXP object. Although the exit can change the user data by calling the `setExitData()` method, the user data is refreshed every time the exit is called. Any changes made to the user data are therefore lost. However, the exit can pass data from one call to the next by using the exit user area of the MQCXP object. The exit accesses the exit user area by reference by calling the `getExitUserArea()` method.

Every exit class must have a constructor. The constructor can be either the default constructor, as shown in the previous example, or a constructor with a string parameter. The constructor is called to create an instance of the exit class for each exit defined in the class. Therefore, in the previous example, an instance of the `MyMQExits` class is created for the send exit, another instance is created for the receive exit, and a third instance is created for the security exit. When a constructor with a string parameter is called, the parameter contains the same user data that is passed to the channel exit for which the instance is being created. If an exit class has both a default constructor and a single parameter constructor, the single parameter constructor takes precedence.

Do not close the connection from within a channel exit.

When data is sent to the server end of a connection, TLS encryption is performed *after* any channel exits are called. Similarly, when data is received from the server end of a connection, TLS decryption is performed *before* any channel exits are called.

In versions of IBM MQ classes for JMS earlier than Version 7.0, channel exits were implemented using the interfaces `MQSendExit`, `MQReceiveExit`, and `MQSecurityExit`. You can still use these interfaces, but the new interfaces are preferred for improved function and performance.

*Configuring IBM MQ classes for JMS to use channel exits:*

An IBM MQ classes for JMS application can use channel security, send, and receive exits on the MQI channel that starts when the application connects to a queue manager. The application can use exits written in Java, C, or C++. The application can also use a sequence of send or receive exits that are run in succession.

The following properties are used specify a send exit, or a sequence of send exits, used by a JMS connection:

- The **SENDEXIT** property of an `MQConnectionFactory` object.
- The **sendexit** property on an activation specification used by the IBM MQ resource adapter for inbound communication,
- The **sendexit** property on a `ConnectionFactory` object used by the IBM MQ resource adapter for output communication.

The value of the property is a string that comprises one or more items separated by commas. Each item identifies a send exit in one of the following ways:

- The name of a class that implements the `WMQSendExit` interface for a send exit written in Java.
- A string in the format *libraryName (entryPointName)* for a send exit written in C or C++.

In a similar way, the following properties specify the receive exit, or sequence of receive exits, used by a connection:

- The **RECEXIT** property of an MQConnectionFactory object.
- The **receiveexit** property on an activation specification used by the IBM MQ resource adapter for inbound communication,
- The **receiveexit** property on a ConnectionFactory object used by the IBM MQ resource adapter for output communication.

The following properties specify the security exit used by a connection:



- The **SECEXIT** property of an MQConnectionFactory object.
- The **securityexit** property on an activation specification used by the IBM MQ resource adapter for inbound communication,
- The **securityexit** property on a ConnectionFactory object used by the IBM MQ resource adapter for output communication.

For MQConnectionFactory, you can set the **SENDEXIT**, **RECEXIT** and **SECEXIT** properties by using the IBM MQ JMS administration tool or IBM MQ Explorer. Alternatively, an application can set the properties by calling the `setSendExit()`, `setReceiveExit()`, and `setSecurityExit()` methods.

Channel exits are loaded by their own class loader. To find a channel exit, the class loader searches the following locations in the specified order.

1. The class path specified by the property **com.ibm.mq.cfg.ClientExitPath.JavaExitsClasspath** or by the **JavaExitsClassPath** attribute in the Channels stanza of the IBM MQ client configuration file.
2. The class path specified by the Java system property **com.ibm.mq.exitClasspath**. Note that this property is now deprecated.
3. The IBM MQ exits directory, as shown in Table 141. The class loader first searches the directory for class files that are not packaged in Java archive (JAR) files. If the channel exit is not found, the class loader then searches the JAR files in the directory.

Table 141. The IBM MQ exits directory

Platform	Directory
 UNIX and Linux	/var/mqm/exits (32-bit channel exits)/var/mqm/exits64 (64-bit channel exits)
 Windows	<i>install_data_dir</i> \exits where <i>install_data_dir</i> is the directory that you chose for the IBM MQ data files during installation. The default directory is C:\ProgramData\IBM\MQ.

**Note:** If a channel exit exists in more than one location, the IBM MQ classes for JMS loads the first instance that it finds.

The parent of the class loader is the class loader that is used to load IBM MQ classes for JMS. It is therefore possible for the parent class loader to load a channel exit if it cannot be found in any of the preceding locations. However, when you are using the IBM MQ classes for JMS in an environment such as a JEE application server, you are not likely to be able to influence the choice of the parent class loader and so the class loader should be configured by setting the Java system property **com.ibm.mq.cfg.ClientExitPath.JavaExitsClasspath** on the application server.

If your application is being run with the Java Security Manager enabled, then the policy configuration file used by the Java runtime environment that the application is running in must have the permissions to load a channel exit class. For information on how to do this, see Running IBM MQ classes for JMS applications under the Java Security Manager.

The MQSendExit, MQReceiveExit, and MQSecurityExit interfaces supplied with versions of IBM WebSphere MQ earlier than Version 7.0 are still supported. If you use channel exits that implement these interfaces, `com.ibm.mq.jar` must be present in the class path.

For information about how to write channel exits in C, see “Channel-exit programs for messaging channels” on page 341. You must store channel exit programs written in C or C++ in the directory shown in Table 141 on page 980.

If your application uses a client channel definition table (CCDT) to connect to a queue manager, see “Using a client channel definition table with IBM MQ classes for JMS” on page 982.

*Specifying the user data to be passed to channel exits when using IBM MQ classes for JMS:*

Up to 32 characters of user data can be passed to a channel exit when it is called.

The `SENDEXITINIT` property of an `MQConnectionFactory` object specifies the user data that is passed to each send exit when it is called. The value of the property is a string that comprises one or more items of user data separated by commas. The position of each item of user data within the string determines which send exit, in a sequence of send exits, the user data is passed to. For example, the first item of user data in the string is passed to the first send exit in a sequence of send exits.

You can set the `SENDEXITINIT` property by using the IBM MQ JMS administration tool or IBM MQ Explorer. Alternatively, an application can set the property by calling the `setSendExitInit()` method.

In a similar way, the `RECEXITINIT` property of a `ConnectionFactory` object specifies the user data that is passed to each receive exit, and the `SECSEXITINIT` property specifies the user data passed to a security exit. You can set these properties by using the IBM MQ JMS administration tool or IBM MQ Explorer. Alternatively, an application can set the properties by calling the `setReceiveExitInit()` and `setSecurityExitInit()` methods.

Note the following rules when specifying user data that is passed to channel exits:

- If the number of items of user data in a string is more than the number of exits in a sequence, the excess items of user data are ignored.
- If the number of items of user data in a string is less than the number of exits in a sequence, each unspecified item of user data is set to an empty string. Two commas in succession within a string, or a comma at the beginning of a string, also denotes an unspecified item of user data.

If an application uses a client channel definition table (CCDT) to connect to a queue manager, any user data specified in a client connection channel definition is passed to channel exits when they are called. For more information about using a client channel definition table, see “Using a client channel definition table with IBM MQ classes for JMS” on page 982.

*Using a client channel definition table with IBM MQ classes for JMS:*

An IBM MQ classes for JMS application can use client connection channel definitions that are stored in a client channel definition table (CCDT). You configure a `ConnectionFactory` object to use the CCDT. There are some restrictions on its use.

As an alternative to creating a client connection channel definition by setting certain properties of a `ConnectionFactory` object, an IBM MQ classes for JMS application can use client connection channel definitions that are stored in a client channel definition table. These definitions are created by IBM MQ Script (MQSC) commands or IBM MQ Programmable Command Format (PCF) commands. When the application creates a `Connection` object, IBM MQ classes for JMS searches the client channel definition table for a suitable client connection channel definition, and uses the channel definition to start an MQI channel. For more information about client channel definition tables and how to construct one, see Client channel definition table.

To use a client channel definition table, the `CCDTURL` property of a `ConnectionFactory` object must be set to a `URL` object. IBM MQ classes for JMS do not read the information about the CCDT from the IBM MQ MQI client configuration file, although some other values are used from there (see “The IBM MQ classes for JMS configuration file” on page 837 for which value apply). The `URL` object encapsulates a uniform resource locator (URL) that identifies the name and location of the file containing the client channel definition table and specifies how the file can be accessed. You can set the `CCDTURL` property by using the IBM MQ JMS administration tool, or an application can set the property by creating a `URL` object and calling the `setCCDTURL()` method of the `ConnectionFactory` object.

For example, if the file `ccdt1.tab` contains a client channel definition table and is stored on the same system on which the application is running, the application can set the `CCDTURL` property in the following way:

```
java.net.URL chanTab1 = new URL("file:///home/admdata/ccdt1.tab");
factory.setCCDTURL(chanTab1);
```

As another example, suppose the file `ccdt2.tab` contains a client channel definition table and is stored on a system that is different from the one on which the application is running. If the file can be accessed using the FTP protocol, the application can set the `CCDTURL` property in the following way:

```
java.net.URL chanTab2 = new URL("ftp://ftp.server/admdata/ccdt2.tab");
factory.setCCDTURL(chanTab2);
```

In addition to setting the `CCDTURL` property of the `ConnectionFactory` object, the `QMANAGER` property of the same object must be set to one of the following values:

- The name of a queue manager
- An asterisk (\*) followed by the name of a queue manager group
- An asterisk (\*)
- An empty string, or a string containing all blank characters

These are the same values that can be used for the `QMgrName` parameter on an `MQCONN` call issued by a client application that is using Message Queue Interface (MQI). For more information about the meaning of these values therefore, see `MQCONN`. You can set the `QMANAGER` property by using the IBM MQ JMS administration tool or IBM MQ Explorer. Alternatively, an application can set the property by calling the `setQueueManager()` method of the `ConnectionFactory` object.

If an application then creates a `Connection` object from the `ConnectionFactory` object, IBM MQ classes for JMS accesses the client channel definition table identified by the `CCDTURL` property, uses the `QMANAGER` property to search the table for a suitable client connection channel definition, and then uses the channel definition to start an MQI channel to a queue manager.

Note that the `CCDTURL` and `CHANNEL` properties of a `ConnectionFactory` object cannot both be set when the application calls the `createConnection()` method. If both properties are set, the method throws an exception. The `CCDTURL` or `CHANNEL` property is considered to be set if its value is anything other than null, an empty string, or a string containing all blank characters.

When IBM MQ classes for JMS finds a suitable client connection channel definition in the client channel definition table, it uses only the information extracted from the table to start an MQI channel. Any channel related properties of the `ConnectionFactory` object are ignored.

In particular, note the following points if you are using TLS:

- An MQI channel uses TLS only if the channel definition extracted from the client channel definition table specifies the name of a CipherSpec supported by IBM MQ classes for JMS.
- A client channel definition table also contains information about the location of Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs). IBM MQ classes for JMS uses only this information to access LDAP servers that hold CRLs.
- A client channel definition table can also contain the location of an OCSP responder. IBM MQ classes for JMS cannot use the OCSP information in a client channel definition table file. However, you can configure OCSP as described in the section Online Certificate Status Protocol (OCSP) in Java and JMS client applications.

For more information about using TLS with a client channel definition table, see [Using the extended transactional client with TLS channels](#).

Note also the following points if you are using channel exits:

- An MQI channel uses only the channel exits and associated user data specified by the channel definition extracted from the client channel definition table.
- A channel definition extracted from a client channel definition table can specify channel exits that are written in Java. This means, for example, that the `SCYEXIT` parameter on the `DEFINE CHANNEL` command to create a client connection channel definition can specify the name of a class that implements the `WMQSecurityExit` interface. Similarly, the `SENDEXIT` parameter can specify the name of a class that implements the `WMQSendExit` interface, and the `RCVEXIT` parameter can specify the name of a class that implements the `WMQReceiveExit` interface. For more information about how to write a channel exit in Java, see [Writing channel exits in Java for IBM MQ classes for JMS](#) on page 978.

The use of channel exits written in a language other than Java is also supported. For information about how to specify the `SCYEXIT`, `SENDEXIT`, and `RCVEXIT` parameters on the `DEFINE CHANNEL` command for channel exits written in another language, see `DEFINE CHANNEL`.

### *Automatic JMS client reconnection:*

Configure your JMS client to reconnect automatically following a network, queue manager, or server failure.

Normally, if a stand-alone IBM MQ classes for JMS application is connected to a queue manager by using the client transport, and the queue manager becomes unavailable for some reason (due to a network outage, a queue manager failure, or the queue manager being stopped, for example), the IBM MQ classes for JMS will throw a `JMSEException` the next time the application tries to communicate with the queue manager. The application must catch the `JMSEException` and attempt to reconnect to the queue manager. You can simplify the design of the application by enabling automatic client reconnection. When the queue manager becomes unavailable, the IBM MQ classes for JMS attempts to reconnect to the queue manager automatically on behalf of the application. This means that the application does not need to contain logic to reconnect.

Automatic client reconnection is only available to stand-alone IBM MQ classes for JMS applications. The use of automatic client reconnection within Java Platform, Enterprise Edition application servers is not supported.

### *Automatic JMS client reconnection by using CONNECTIONNAMELIST:*

If a stand-alone IBM MQ classes for JMS application uses a Connection Factory that has the `CONNECTIONNAMELIST` property set, the application is eligible to use automatic client reconnection.

The behavior of the automatic client reconnection functionality that is provided by the IBM MQ classes for JMS depends on the properties that follow:

#### **The JMS Connection Factory property TRANSPORT (Short name TRAN)**

`TRANSPORT` specifies how applications that use the Connection Factory connects to a queue manager. This property must be set to the value `CLIENT` for automatic client reconnection to be used. Automatic client reconnection is not available to applications that connect to a queue manager that uses a Connection Factory that has the `TRANSPORT` property set to `BIND`, `DIRECT`, or `DIRECTHTTP`.

#### **The JMS Connection Factory property QMANAGER (Short name QMGR)**

The `QMANAGER` property specifies the name of the queue manager that the Connection Factory connects to.

#### **The JMS Connection Factory property CONNECTIONNAMELIST (Short name CRHOSTS)**

The `CONNECTIONNAMELIST` property is a comma-separated list, where each entry contains information on the host name and port that are to be used to connect to the queue manager specified by the `QMANAGER` property when you are using the `CLIENT` transport. The list has the following format: host name(port), host name(port).

#### **The JMS Connection Factory property CLIENTRECONNECTOPTIONS (Short name CROPT)**

`CLIENTRECONNECTOPTIONS` controls whether the IBM MQ classes for JMS will attempt to automatically connect to a queue manager on behalf of an application if a queue manager becomes available.

#### **The DefRecon attribute in the Channels stanza of the client configuration file**

The `DefRecon` attribute provides an administrative option to enable all applications to automatically reconnect, or to disable the automatic reconnection for applications that are written to reconnect automatically.

Automatic client reconnection is only available when an application successfully connects to a queue manager.



When an application connects to a queue manager that uses the CLIENT transport, the IBM MQ classes for JMS use the value of the Connection Factory property CLIENTRECONNECTOPTIONS to determine whether to use automatic client reconnection, if the queue manager that the application is connected to becomes unavailable. Table 1 shows the possible values for the CLIENTRECONNECTOPTIONS property, and the behavior of the IBM MQ classes for JMS for each of these values:

Table 142. Possible CLIENTRECONNECTOPTIONS property values.

CLIENTRECONNECTOPTIONS	Behavior of IBM MQ classes for JMS
ANY	Use the value of the CONNECTIONNAMELIST property to open a connection to a host name and port combination, and connect to any queue manager. In order to use this automatic client reconnection option, the QMANAGER property must be set to either the default value or "*".
ASDEF	Use the value of DefRecon to determine whether automatic client reconnection is available.
DISABLED	Do not perform any automatic client reconnection and return a JMSEException to the application.
QMGR	Use the value of the CONNECTIONNAMELIST property to open a connection to a host name and port combination, and connect to the queue manager specified by the QMANAGER property.

When you perform automatic client reconnection, the IBM MQ classes for JMS uses the information in the Connection Factory property CONNECTIONNAMELIST to determine what system to reconnect to.

The IBM MQ classes for JMS initially tries to reconnect by using the host name and port that is specified in the first entry in the CONNECTIONNAMELIST. If a connection is made, the IBM MQ classes for JMS then tries to connect to the queue manager that has the name specified in the QMANAGER property. If a connection to the queue manager can be established, the IBM MQ classes for JMS reopens all of the IBM MQ objects that the application had open before automatic client reconnection and continue running as before.

If a connection cannot be established to the required queue manager by using the first entry in the CONNECTIONNAMELIST, the IBM MQ classes for JMS tries the second entry in the CONNECTIONNAMELIST, and so on.

When the IBM MQ classes for JMS have tried all of the entries in the CONNECTIONNAMELIST, they wait for a period of time before they try to reconnect again. To perform the new reconnection attempt, the IBM MQ classes for JMS start with the first entry in the CONNECTIONNAMELIST. They then try each entry in the CONNECTIONNAMELIST in turn until either a reconnection occurs or the end of the CONNECTIONNAMELIST is reached, at which point the IBM MQ classes for JMS waits for a period of time before they try again.

This process of automatic client reconnection continues until the IBM MQ classes for JMS successfully reconnected to the queue manager specified by the QMANAGER property.

By default, the reconnection attempts happen at the following intervals:

- The first attempt is made after an initial delay of 1 second, plus a random element up to 250 milliseconds.
- The second attempt is made 2 seconds, plus a random interval of up to 500 milliseconds, after the first attempt fails.
- The third attempt is made 4 seconds, plus a random interval of up to 1 second, after the second attempt fails.

- The fourth attempt is made 8 seconds, plus a random interval of up to 2 seconds, after the third attempt fails.
- The fifth attempt is made 16 seconds, plus a random interval of up to 4 seconds, after the fourth attempt fails.
- The sixth attempt, and all subsequent attempts are made 25 seconds, plus a random interval of up to 6 seconds and 250 milliseconds after the previous attempt fails.

The reconnection attempts are delayed by intervals that are partly fixed and partly random. This is to prevent all of the IBM MQ classes for JMS applications that were connected to a queue manager that is no longer available from reconnecting simultaneously.

If you need to increase the default values, to more accurately reflect the amount of time that is required for a queue manager to recover, or a standby queue manager to become active, modify the ReconDelay attribute in the Channel stanza of the client configuration file, for more information, see CHANNELS stanza of the client configuration file.

Whether an IBM MQ classes for JMS application continues to work correctly after being reconnected automatically depends on its design. Read the related topics to understand how to design applications can use the automatic reconnection functionality.

*Connecting to multi-instance queue managers using CONNECTIONNAMELIST:*

Automatic client reconnection can be used by the IBM MQ classes for JMS applications that connect to a multi-instance queue manager.

If the queue manager instance that an application is using becomes unavailable, the IBM MQ classes for JMS can automatically try to connect to the standby instance on behalf of the application. The application blocks while the automatic client reconnection is taking place, and resumes when the IBM MQ classes for JMS establish a connection to the standby queue manager.

To enable automatic client reconnection for a multi-instance queue manager, set the following properties on the Connection Factory that is used by the IBM MQ classes for JMS application:

**CHANNEL**

The name of a server connection channel defined on the queue manager.

**QMANAGER**

The name of the multi-instance queue manager.

**CONNECTIONNAMELIST=host1(port1), host2(port2).**

The first entry in the list must contain the host name and the port that is used to contact the primary queue manager instance. The second entry should contain the host name, and the port of the system where the standby queue manager instance is located.

**CLIENTRECONNECTOPTIONS=QMGR.**

This ensures that the IBM MQ classes for JMS tries to reconnect to a queue manager that has the same name as the queue manager that the application was previously connected to.

*Automatic JMS client reconnection with CCDTs:*

If a stand-alone IBM MQ classes for JMS application uses a Connection Factory that has the CCDTURL property set, the application is eligible to use automatic client reconnection.

The behavior of the automatic client reconnection functionality that is provided by the IBM MQ classes for JMS depends on the following properties:

**The JMS Connection Factory property TRANSPORT (Short name TRAN)**

TRANSPORT specifies how applications that use the Connection Factory connects to a queue manager. This property must be set to the value CLIENT for automatic client reconnection to be used. Automatic client reconnection is not available to applications that connect to a queue manager by using a Connection Factory that has the TRANSPORT property set to BIND, DIRECT, or DIRECTHTTP.

**The JMS Connection Factory property QMANAGER (Short name QMGR)**

The QMANAGER property specifies the name of the queue manager that the Connection Factory connects to.

**The JMS Connection Factory property CCDTURL (Short name CCDT)**

The CCDTURL property points to the client channel definition table that the IBM MQ classes for JMS uses when it connects to a queue manager.

**The JMS Connection Factory property CLIENTRECONNECTOPTIONS (Short name CROPT)**

CLIENTRECONNECTOPTIONS controls whether the IBM MQ classes for JMS attempts to automatically connect to a queue manager on behalf of an application if a queue manager becomes available.

**The DefRecon attribute in the Channels stanza of the client configuration file**

The DefRecon attribute provides an administrative option to enable all applications to automatically reconnect, or to disable the automatic reconnection for applications that are written to reconnect automatically.

Automatic client reconnection is only available when an application successfully connects to a queue manager.

When an application connects to a queue manager by using the CLIENT transport, the IBM MQ classes for JMS use the value of the Connection Factory property CLIENTRECONNECTOPTIONS to determine whether to use automatic client reconnection, if the queue manager that the application is connected to becomes unavailable. Table 1 shows the possible values for the CLIENTRECONNECTOPTIONS property, and the behavior of the IBM MQ classes for JMS for each of these values:

*Table 143. Possible CLIENTRECONNECTOPTIONS property*

CLIENTRECONNECTOPTIONS	Behavior of IBM MQ classes for JMS
ANY	Open the client channel definition table that is specified by the CCDTURL property, pick an entry in the table, and then use that entry to start a client connection channel to a queue manager. To use this automatic client reconnection option, the QMANAGER property must be set to either: <ul style="list-style-type: none"><li>• An asterisk (*)</li><li>• An asterisk (*) followed by the name of a queue manager group</li><li>• An empty string, or a string that contains all blank characters</li></ul>

Table 143. Possible CLIENTRECCECTOPTIONS property (continued)

CLIENTRECONNECTOPTIONS	Behavior of IBM MQ classes for JMS
ASDEF	Use the value of DefRecon to determine whether automatic client reconnection is available.
DISABLED	Do not perform any automatic client reconnection and return a JMSEException to the application.
QMGR	Open the client channel definition table that is specified by the CCDTURL property, find the entries in the table that match the queue manager name that is specified by the QMANAGER property and then use those entries to start a client connection channel to that queue manager.

When performing automatic client reconnection, the IBM MQ classes for JMS uses the client channel definition table that is specified in the CCDTURL property to determine what system to reconnect to.

The IBM MQ classes for JMS initially parses the client channel definition table and finds a suitable entry that matches the value of the QMANAGER property. When an entry is found, the IBM MQ classes for JMS tries to reconnect to the required queue manager using that entry. If a connection to the queue manager can be established, the IBM MQ classes for JMS reopens all of the IBM MQ objects that the application had open before automatic client reconnection and continue running as before.

If a connection cannot be established to the required queue manager, the IBM MQ classes for JMS looks for another suitable entry in the client channel definition table and tries to use that, and so on.

When the IBM MQ classes for JMS have tried all of the suitable entries in the client channel definition table, they wait for a period of time before trying to reconnect again. To perform the new reconnection attempt, the IBM MQ classes for JMS parses the client channel definition table again and tries the first suitable entry. They will then try each suitable entry in the client channel definition table in turn until either a reconnection occurs or the last suitable entry in the client channel definition table has been tried, at which point the IBM MQ classes for JMS waits for a period of time before trying again.

This process of automatic client reconnection continues until the IBM MQ classes for JMS successfully reconnects to the queue manager specified by the QMANAGER property.

By default, the reconnection attempts happen at the following intervals:

- The first attempt is made after an initial delay of 1 second, plus a random element up to 250 milliseconds.
- The second attempt is made 2 seconds, plus a random interval of up to 500 milliseconds, after the first attempt fails.
- The third attempt is made 4 seconds, plus a random interval of up to 1 second, after the second attempt fails.
- The fourth attempt is made 8 seconds, plus a random interval of up to 2 seconds, after the third attempt fails.
- The fifth attempt is made 16 seconds, plus a random interval of up to 4 seconds, after the fourth attempt fails.
- The sixth attempt, and all subsequent attempts are made 25 seconds, plus a random interval of up to 6 seconds and 250 milliseconds after the previous attempt fails.

The reconnection attempts are delayed by intervals that are partly fixed and partly random. This prevents all of the IBM MQ classes for JMS applications, that were connected to a queue manager that is no longer available, from reconnecting simultaneously.

If you need to increase the default values, to more accurately reflect the amount of time that is required for a queue manager to recover, or a standby queue manager to become active, modify the ReconDelay attribute in the Channel stanza of the client configuration file, for more information, see CHANNELS stanza of the client configuration file.

Whether an IBM MQ classes for JMS application continues to work correctly after it has been reconnected automatically depends on its design. Read the related topics to understand how to design applications that can use the automatic reconnection functionality.

*Connecting to multi-instance queue managers by using CCDTs:*

Automatic client reconnection can be used by IBM MQ classes for JMS applications that connect to a multi-instance queue manager.

If the queue manager instance that an application is using becomes unavailable, the IBM MQ classes for JMS can automatically try to connect to the standby instance on behalf of the application. The application blocks while the automatic client reconnection is taking place, and resumes when the IBM MQ classes for JMS establishes a connection to the standby queue manager.

To enable automatic client reconnection for a multi-instance queue manager, set the following properties on the Connection Factory that is used by the IBM MQ classes for JMS application:

**QMANAGER=The name of the multi-instance queue manager.**

**CCDTURL=URI**

The URI for a client channel definition table that contains two entries for the multi-instance queue manager; one for the primary instance and one for the stand-by instance.

*Using automatic client reconnection in Java Platform, Enterprise Edition environments:*

Automatic client reconnection with activation specifications is not supported.

The IBM MQ resource adapter provides its own mechanism for reconnecting activation specifications if the queue manager that the activation specification was connecting to becomes unavailable.

This mechanism is controlled by:

- The IBM MQ resource adapter property `reconnectionRetryCount`.
- The IBM MQ resource adapter property `reconnectionRetryInterval`.
- The activation specification property `connectionNameList`.

For more information on these properties, see “Configuration for ResourceAdapter object properties” on page 1104.

The use of automatic client reconnection within a message-driven bean application's `onMessage()` method, or any other application that is running within the Java Platform, Enterprise Edition environment is not supported. The application needs to implement its own reconnection logic if the queue manager it was connecting to becomes unavailable.

*Sharing a TCP/IP connection in IBM MQ classes for JMS:*

Multiple instances of an MQI channel can be made to share a single TCP/IP connection.

Applications that are running inside the same Java runtime environment, and that use the IBM MQ classes for JMS or the IBM MQ resource adapter to connect to a queue manager by using the CLIENT transport, can be made to share the same channel instance.

There is a one-to-one relationship between channel instances and TCP/IP connections. One TCP/IP connection is created for every channel instance.

If a channel is defined with the **SHARECNV** parameter set to a value greater than 1, then that number of conversations can share a channel instance. To enable a connection factory or an activation specification to use this function, set the **SHARECONVALLOWED** property to YES.

Every JMS connection and JMS session that is created by a JMS application creates its own conversation with the queue manager.

When an activation specification starts up, the IBM MQ resource adapter starts a conversation with the queue manager for the activation specification to use. Every server session in the server session pool that is associated with the activation specification also starts a conversation with the queue manager.

The **SHARECNV** attribute is a best effort approach to connection sharing. Therefore, when a **SHARECNV** value greater than 0 is used with the IBM MQ classes for JMS, it is not guaranteed that a new connection request will always share an already established connection.

### Calculating the number of channel instances

Use the following formulae to determine the maximum number of channel instances that are created by an application:

#### Activation specifications

Number of channel instances =  $(maxPoolDepth\_value + 1) / SHARECNV\_value$

Where *maxPoolDepth\_value* is the value of the **maxPoolDepth** property and *SHARECNV\_value* is the value of the **SHARECNV** property on the channel that is used by the activation specification.

#### Other JMS applications

Number of channel instances =  $(jms\_connections + jms\_sessions) / SHARECNV\_value$

Where *jms\_connections* is the number of connections that are created by the application, *jms\_sessions* is the number of JMS sessions that are created by the application, and *SHARECNV\_value* is the value of the **SHARECNV** property on the channel that is used by the activation specification.

### Examples

The following examples show how to use the formulae to calculate the number of channel instances that are created on a queue manager by applications by using either the IBM MQ classes for JMS or the IBM MQ resource adapter.

#### JMS application example

A JMS application connection connects to a queue manager by using the CLIENT transport, and creates a JMS connection and three JMS sessions. The channel that the application is using to connect to the queue manager has the **SHARECNV** property set to the value of 10. When the application is running, there are four conversations between the application and the queue manager and one channel instance. The four conversations all share the channel instance.

### Activation specification example

An activation specification connects to a queue manager by using the CLIENT transport. The activation specification is configured with the **maxPoolDepth** property set to 10. The channel that the activation specification is configured to use has the **SHARECNV** property set to 10. When the activation specification is running, and processing 10 messages concurrently, the number of conversations between the activation specification and the queue manager is 11 (10 conversations for the server sessions, and one for the activation specification). The number of channel instances that are used by the activation specification is 2.

### Activation specification example

An activation specification connects to a queue manager by using the CLIENT transport. The activation specification is configured with the **maxPoolDepth** property set to 5. The channel that the activation specification is configured to use has the **SHARECNV** property set to 0. When the activation specification is running, and processing 5 messages concurrently, the number of conversations between the activation specification and the queue manager is 6 (five conversations for the server sessions, and one for the activation specification). The number of channel instances that are used by the activation specification is 6, because the **SHARECNV** property on the channel is set to 0, every conversation uses its own channel instance.

*Specifying a range of ports for client connections in IBM MQ classes for JMS:*

Use the LOCALADDRESS property to specify a range of ports that your application can bind to.

When an IBM MQ classes for JMS application attempts to connect to an IBM MQ queue manager in client mode, a firewall might allow only those connections that originate from specified ports or a range of ports. In this situation, you can use the LOCALADDRESS property of a ConnectionFactory, QueueConnectionFactory, or TopicConnectionFactory object to specify a port, or a range of ports, that the application can bind to.

You can set the LOCALADDRESS property by using the IBM MQ JMS administration tool, or by calling the setLocalAddress() method in a JMS application. Here is an example of setting the property from within an application:

```
mqConnectionFactory.setLocalAddress("192.0.2.0(2000,3000)");
```

When the application connects to a queue manager subsequently, the application binds to a local IP address and port number in the range 192.0.2.0(2000) to 192.0.2.0(3000).

In a system with more than one network interface, you can also use the LOCALADDRESS property to specify which network interface must be used for a connection.

For a real-time connection to a broker, the LOCALADDRESS property is relevant only when multicast is used. In this case, you can use the property to specify which local network interface must be used for a connection, but the value of the property must not contain a port number, or a range of port numbers.

Connection errors might occur if you restrict the range of ports. If an error occurs, a JMSEException is thrown with an embedded MQException that contains the IBM MQ reason code MQRC\_Q\_MGR\_NOT\_AVAILABLE and the following message:

```
Socket connection attempt refused due to LOCAL_ADDRESS_PROPERTY restrictions
```

An error might occur if all the ports in the specified range are in use, or if the specified IP address, host name, or port number is not valid (a negative port number, for example).

Because IBM MQ classes for JMS might create connections other than those required by an application, always consider specifying a range of ports. In general, every session created by an application requires one port and IBM MQ classes for JMS might require three or four additional ports. If a connection error does occur, increase the range of ports.

Connection pooling, which is used by default in IBM MQ classes for JMS, might have an effect on the speed at which ports can be reused. As a result, a connection error might occur while ports are being freed.

*Channel compression in IBM MQ classes for JMS:*

An IBM MQ classes for JMS application can use IBM MQ facilities to compress a message header or data.

Compressing the data that flows on an IBM MQ channel can improve the performance of the channel and reduce network traffic. Using function supplied with IBM MQ, you can compress the data that flows on message channels and MQI channels. On either type of channel, you can compress header data and message data independently of each other. By default, no data is compressed on a channel.

An IBM MQ classes for JMS application specifies the techniques that can be used for compressing header or message data on a connection by creating a `java.util.Collection` object. Each compression technique is an `Integer` object in the collection, and the order in which the application adds the compression techniques to the collection is the order in which the compression techniques are negotiated with the queue manager when the application creates the connection. The application can then pass the collection to a `ConnectionFactory` object by calling the `setHdrCompList()` method, for header data, or the `setMsgCompList()` method, for message data. When the application is ready, it can create the connection.

The following code fragments illustrate the approach described. The first code fragment shows you how to implement header data compression:

```
Collection headerComp = new Vector();
headerComp.add(new Integer(WMQConstants.WMQ_COMPHDR_SYSTEM));
.
.
.
((MQConnectionFactory) cf).setHdrCompList(headerComp);
.
.
connection = cf.createConnection();
```

The second code fragment shows you how to implement message data compression:

```
Collection msgComp = new Vector();
msgComp.add(new Integer(WMQConstants.WMQ_COMPMSG_RLE));
msgComp.add(new Integer(WMQConstants.WMQ_COMPMSG_ZLIB_HIGH));
.
.
.
((MQConnectionFactory) cf).setMsgCompList(msgComp);
.
.
connection = cf.createConnection();
```

In the second example, the compression techniques are negotiated in the order RLE, then ZLIB\_HIGH, when the connection is created. The compression technique that is selected cannot be changed during the lifetime of the `Connection` object. To use compression on a connection, the `setHdrCompList()` and the `setMsgCompList()` methods must be called before creating the `Connection` object.



*Putting messages asynchronously in IBM MQ classes for JMS:*

Normally, when an application sends messages to a destination, the application has to wait for the queue manager to confirm that it has processed the request. You can improve messaging performance in some circumstances by choosing instead to put messages asynchronously. When an application puts a message asynchronously, the queue manager does not return the success or failure of each call, but you can instead check for errors periodically.

Whether a destination returns control to the application, without determining whether the queue manager has received the message safely, depends upon the following properties:

- The JMS Destination Property `PUTASYNCALLOWED` (short name - `PAALD`).  
`PUTASYNCALLOWED` controls whether JMS applications can put messages asynchronously, if the underlying queue or topic that the JMS Destination represents, allows this option.
- The IBM MQ queue or topic property `DEFPRESP` (Default put response type).  
`DEFPRESP` specifies whether applications that put messages to the queue, or publish messages to the topic, can make use of the asynchronous put functionality.

The following table shows the possible values for the `PUTASYNCALLOWED` and `DEFPRESP` properties, and what values are required for the asynchronous put functionality to be enabled:

*Table 144. PUTASYNCALLOWED and DEFPRESP properties determining if messages are put asynchronously.*

JMS Destination property	PUTASYNCALLOWED = NO	PUTASYNCALLOWED = YES	PUTASYNCALLOWED = AS_DEST or AS_Q_DEF or AS_T_DEF
IBM MQ queue property			
DEFPRESP=SYNC	Asynchronous put functionality not enabled	Asynchronous put functionality enabled	Asynchronous put functionality not enabled
DEFPRESP=ASYNC	Asynchronous put functionality not enabled	Asynchronous put functionality enabled	Asynchronous put functionality enabled

For messages sent in a transacted session, the application ultimately determines whether the queue manager has received the messages safely when it calls `commit()`.

If an application sends persistent messages within a transacted session, and one or more of the messages are not received safely, the transaction fails to commit and produces an exception. However, if an application sends nonpersistent messages within a transacted session, and one or more of the messages are not received safely, the transaction commits successfully. The application does not receive any feedback that the nonpersistent messages did not arrive safely.

For nonpersistent messages sent in a session that is not transacted, the `SENDCHECKCOUNT` property of the *ConnectionFactory* object specifies how many messages are to be sent, before IBM MQ classes for JMS checks that the queue manager has received the messages safely.

If a check discovers that one or more messages were not received safely, and the application has registered an exception listener with the connection, IBM MQ classes for JMS calls the `onException()` method of the exception listener to pass a JMS exception to the application.

The JMS exception has an error code of `JMSWMQ0028` and this code displays the following message:  
At least one asynchronous put message failed or gave a warning.

The JMS exception also has a linked exception that provides more details. The default value of the `SENDCHECKCOUNT` property is zero, which means that no such checks are made.

This optimization is of most benefit to an application that connects to a queue manager in client mode, and needs to send a sequence of messages in rapid succession, but does not require immediate feedback from the queue manager for each message sent. However, an application can still use this optimization even if it connects to a queue manager in bindings mode, but the expected performance benefit is not as great.

*Using read ahead with IBM MQ classes for JMS:*

The read ahead functionality that is provided by IBM MQ allows non-persistent messages that are received outside of a transaction to be sent to the IBM MQ classes for JMS before an application requests them. The IBM MQ classes for JMS store the messages in an internal buffer, and pass the messages to the application when the application asks for them.

IBM MQ classes for JMS applications that use `MessageConsumers` or `MessageListeners` to receive messages from a destination outside of a transaction can use the read ahead functionality. Using read ahead allows applications that use these objects to benefit from improved performance when they receive messages.

Whether an application that uses `MessageConsumers` or `MessageListeners` can use read ahead depends upon the following properties:

- The JMS Destination Property `READAHEADALLOWED` (short name - `RAALD` ). `READAHEADALLOWED` controls whether JMS applications can use read ahead when getting or browsing non-persistent messages outside of a transaction, if the underlying queue or topic that the JMS Destination represents, allows this option.
- The IBM MQ queue or topic property `DEFREADA` (Default read ahead). `DEFREADA` specifies whether applications that are receiving or browsing non-persistent messages outside of a transaction can use read ahead.

The following table shows the possible values for the `READAHEADALLOWED` and `DEFREADA` properties, and what values are required for the read ahead functionality to be enabled:

*Table 145. READAHEADALLOWED and DEFREADA properties determining if read ahead is used when receiving or browsing non-persistent messages outside of a transaction.*

IBM MQ destination property	READAHEADALLOWED= YES	READAHEADALLOWED= NO	AS_DEST or AS_Q_DEF or AS_T_DEF
IBM MQ queue property			
DEFREADA = NO	Read ahead functionality enabled	Read ahead functionality not enabled	Read ahead functionality not enabled
DEFREADA = YES	Read ahead functionality enabled	Read ahead functionality not enabled	Read ahead functionality enabled
DEFREADA = DISABLED	Read ahead functionality not enabled	Read ahead functionality not enabled	Read ahead functionality not enabled

If the read ahead functionality is enabled, when a `MessageConsumer` or `MessageListener` is created by an application, the IBM MQ classes for JMS create an internal buffer for the destination that the `MessageConsumer` or `MessageListener` is monitoring. There is one internal buffer for each `MessageConsumer` or `MessageListener`. The queue manager starts sending non-persistent messages to the IBM MQ classes for JMS when the application calls one of the following methods:

- `MessageConsumer.receive()`
- `MessageConsumer.receive(long timeout)`
- `MessageConsumer.receiveNoWait()`
- `Session.setMessageListener(MessageListener listener)`

The IBM MQ classes for JMS automatically returns the first message back to the application, by the method call that the application has made. The other non-persistent messages are stored by the IBM MQ classes for JMS in the internal buffer that was created for the destination. When the application requests the next message to process, the IBM MQ classes for JMS will return the next message in the internal buffer.

The IBM MQ classes for JMS requests more non-persistent messages from the queue manager when the internal buffer is empty.

The internal buffer that is used by the IBM MQ classes for JMS is deleted when an application closes a `MessageConsumer`, or the JMS Session that a `MessageListener` is associated with.

For `MessageConsumers`, any unprocessed messages in the internal buffer is lost.

When using `MessageListeners`, what happens to the messages in the internal buffer depends upon the JMS destination property `READAHEADCLOSEPOLICY` (short name - `RACP`). The default value of the property is `DELIVER_ALL`, which means that the JMS session that was used to create the `MessageListener` is not closed until all of the messages in the internal buffer are delivered to the application. If the property is set to `DELIVER_CURRENT`, then the JMS session will be closed after the current message has been processed by the application and all of the remaining messages in the internal buffer are discarded.

*Retained publications in IBM MQ classes for JMS:*

An IBM MQ classes for JMS client can be configured to use retained publications.

A publisher can specify that a copy of a publication must be retained so that it can be sent to future subscribers who register an interest in the topic. This is done in IBM MQ classes for JMS by setting the integer property `JMS_IBM_RETAIN` to the value 1. Constants have been defined for these values in the `com.ibm.msg.client.jms.JmsConstants` interface. For example, if you have created a message `msg`, to set it as a retained publication use the following code:

```
// set as a retained publication
msg.setIntProperty(JmsConstants.JMS_IBM_RETAIN, JmsConstants.RETAIN_PUBLICATION);
```

You can now send the message as normal. `JMS_IBM_RETAIN` can also be queried in a received message. It is therefore possible to query whether a received message is a retained publication.

### **XA support in IBM MQ classes for JMS:**

JMS supports XA-compliant transactions in bindings and client modes with a supported transaction manager within a JEE container.

If you require XA functionality in an application server environment, you must configure your application appropriately. Refer to your application server's own documentation for information about how to configure applications to use distributed transactions.

An IBM MQ queue manager cannot act as a transaction manager for JMS.

## Using JMS 2.0 functionality:

JMS 2.0 introduces several new areas of functionality to the IBM MQ classes for JMS.

When you are developing a JMS application for IBM MQ Version 8.0 or later, you might need to consider the impact of this functionality on your queue manager.

### Related information:

IBM MQ Java language interfaces

#### *JMS 2.0 delivery delay:*

With JMS 2.0, you can specify a delivery delay when sending a message. The queue manager does not deliver the message until after the specified delivery delay has elapsed.

An application can specify a delivery delay in milliseconds, when it sends a message, by using either `MessageProducer.setDeliveryDelay(long deliveryDelay)` or `JMSProducer.setDeliveryDelay(long deliveryDelay)`. This value is added to the time at which the message is sent and gives the earliest time at which any other application can get that message.

In IBM MQ Version 8.0 and later, delivery delay is implemented by using a single internal staging queue. Messages that have a nonzero delivery delay are placed on this queue with a header that indicates the delivery delay and information about the target queue. A component of the queue manager that is called the delivery delay processor monitors the messages on the staging queue. When a message's delivery delay completes, the message is taken off the staging queue and placed on the target queue.

## Messaging clients

The IBM MQ implementation of delivery delay is only available for use when you are using the JMS client. The following restrictions apply if you are using delivery delay with IBM MQ. These restrictions apply equally to `MessageProducers` and `JMSProducers`, but `JMSRuntimeExceptions` are thrown in the case of `JMSProducers`.

- Any attempt to call `MessageProducer.setDeliveryDelay` with a nonzero value when connected to an IBM MQ queue manager earlier than Version 8.0, results in a `JMSException` with a `MQRC_FUNCTION_NOT_SUPPORTED` message.
- Delivery delay is not supported for clustered destinations that have a **DEFBIND** value other than `MQBND_BIND_NOT_FIXED`. If a `MessageProducer` has a nonzero delivery delay set and an attempt is made to send to a destination that does not meet this requirement, then the call results in a `JMSException` with an `MQRC_OPTIONS_ERROR` message.
- Any attempt to set a time to live value that is less than a previously specified nonzero delivery delay, or vice versa, results in a `JMSException` with a `MQRC_EXPIRY_ERROR` message. This checking is done on calling `setTimeToLive` or `setDeliveryDelay` or `send` methods, depending on the exact set of operations chosen.
- Use of retained publications and delivery delay is not supported. Attempting to publish a message with a delivery delay if that message has been marked as retained by using `msg.setIntProperty(JmsConstants.JMS_IBM_RETAIN, JmsConstants.RETAIN_PUBLICATION)` results in a `JMSException` with a `MQRC_OPTIONS_ERROR` message.
- Delivery delay and message grouping is not supported and any attempt to use this combination results in a `JMSException` with a `MQRC_OPTIONS_ERROR` message.

Any failure to send a message with delivery delay results in the client throwing a `JMSException` with a suitable error message, for example queue full. In some situations, the error message might apply to the target destination, or the staging queue, or both.

**Note:** IBM MQ allows applications that put a message in a unit of work to get the same message again even though the unit of work has not committed. This technique does not work with delivery delay as the message is not placed on the staging queue until the unit of work is committed, and as a result will not have been sent to the target destination.

## Authorization

IBM MQ carries out authorization checks on the original target destination when the application sends a message with a nonzero delivery delay. If the application is not authorized, then the send fails. When the queue manager detects that a message's delivery delay is complete, it opens the target queue. No authorization checks are carried out at this point.

## SYSTEM.DDELAY.LOCAL.QUEUE

A new system queue, SYSTEM.DDELAY.LOCAL.QUEUE, is used to implement delivery delay.

- **Multi** On Multiplatforms, SYSTEM.DDELAY.LOCAL.QUEUE exists by default. The system queue must be altered so that its MAXMSGL and MAXDEPTH attributes are sufficient for the expected load.
- **z/OS** On IBM MQ for z/OS, SYSTEM.DDELAY.LOCAL.QUEUE is used as a staging queue for messages that are sent with delivery delay to both local and shared queues. On z/OS, the queue must be created and must either be defined so that its MAXMSGL and MAXDEPTH attributes are sufficient for the expected load.

When this queue is created, it must be secured so that as few users as possible have access to it. Access to the queue must be for maintenance and monitoring purposes only.

When a message is sent by a JMS application with a nonzero delivery delay, it is put to this queue with a new message ID. The original message ID is placed in the correlation ID of the message. This correlation ID allows an application to retrieve a message from the staging queue when required, for example if a large delivery delay was used by mistake.

### **z/OS**

## Considerations for z/OS

If your system is running on z/OS, there are additional considerations to take into account if you want to use delivery delay.

If delivery delay is to be used, the system queue SYSTEM.DDELAY.LOCAL.QUEUE must be defined. It must be defined with a storage class that is sufficient for its expected load, and with HARDENBO, STGCLASS('DEFAULT'), INDXTYPE(NONE), and MSGDLVSQ(FIFO) specified. A sample definition of the system queue is provided, commented out, in the CSQ4INSG JCL.

Delivery delay is not protected by OPMODE. If you use delivery delay with an IBM MQ Version 8.0 queue manager and then migrate back to an earlier release any messages on the SYSTEM.DDELAY.LOCAL.QUEUE queues are marooned unless you manually deal with them.

## Shared queues

Delivery delay is supported for sending messages to shared queues. However, there is only a single, private staging queue that is used regardless of whether the target queue is shared or not. The queue manager that owns that private queue must be running to send the delayed message to its target shared queue when the delay completes.

**Note:** If a non-persistent message is put with a delivery delay to a shared queue, and the queue manager that owns the staging queue shuts down, the original message is lost. As a result non-persistent messages sent with delivery delay to a shared queue are more likely to be lost than non-persistent messages sent without delivery delay to a shared queue.

### Target destination resolution

If the message is sent to a queue, resolution is driven twice; once by the JMS application and once by the queue manager when it takes the message off the staging queue and sends it to the target queue.

Target subscriptions for publications are matched when the JMS application calls the send method.

If a message is sent with persistence or priority according to the queue definition, then the value is set on the first resolution and not the second.

### Expiry interval

Delivery delay preserves the behavior of the expiry property, MQMD.Expiry. For example, if a message was put from a JMS application with an expiry interval of 20,000 ms and a delivery delay of 5,000 ms, and got after an elapsed time of 10,000 ms, then the value of the MQMD.expiry field might be approximately 50 tenths of a second. This value indicates that 15 seconds has elapsed from the time the message was put, to the time when it was got.

If a message expires while on the staging queue and one of the MQRO\_EXPIRATION\_\* options is set, then the report generated is for the original message as sent by the application, the header used to contain the delivery delay information is removed.

### Stopping and starting the delivery delay processor

**z/OS** On z/OS, the delivery delay processor is integrated into the queue manager MSTR address space. When the queue manager starts, the delivery delay processor also starts. If the staging queue is available, it opens the queue and waits for messages to arrive on it to be processed. If the staging queue has not been defined, or is disabled for gets, or another error occurs, the delivery delay processor shuts down. If the staging queue is later defined, or altered to be get enabled, the delivery delay processor restarts. If the delivery delay processor shuts down for any other reason, it can be restarted by altering the PUT attribute of the staging queue from ENABLED to DISABLED and back to ENABLED again. Should you need to stop the delivery delay processor for any reason, set the PUT attribute of the staging queue to DISABLED.

**Multi** On Multiplatforms, the delay processor starts with the queue manager, and is automatically restarted in the event of a recoverable failure.

### Failure to put to target queue

If a delayed message cannot be put to the target queue once its delay completes, the message is dealt with as indicated in its report options: it is either discarded or sent to the dead letter queue. If this action fails, then an attempt is made to put the message later. If the action is successful an exception report is generated and sent to the specified queue, if the report is requested. If the report message could not be sent, the report message is sent to the dead letter queue. If sending the report to the dead letter queue fails and the message is persistent, all changes are discarded and the original message rolled back and redelivered later. If the message is non-persistent the report message is discarded, but other changes are committed. If a delayed publication cannot be delivered because a subscriber has unsubscribed, or in the case of a non-durable subscriber, because it has disconnected, then the message is discarded silently. Report messages are still generated as described earlier.

If a delayed publication cannot be delivered to a subscriber and is instead put to the dead letter queue, and the put to the dead letter queue fails then the message is discarded.

To reduce the likelihood of the put to the target queue failing after the delivery delay has completed, the queue manager performs some basic checks when the JMS client sends a message with a nonzero delivery delay. These checks include whether the queue is put disabled, if the message is bigger than the maximum message length allowed, and if the queue is full.

### **Publish/subscribe**

Matching of a publication to available subscriptions occurs when the JMS application sends a message with a nonzero delivery delay. A message for each matching subscriber is put to the `SYSTEM.DDELAY.LOCAL.QUEUE` queue, where it is kept until the delivery delay completes. If one of those subscribers is a proxy subscription for another queue manager, then the fan-out on that queue manager occurs after the delivery delay is complete. This might result in subscribers on the other queue manager receiving publications that were originally published before they subscribed. This is a deviation from the JMS 2.0 specification.

Delivery delay with publish/subscribe is only supported if the target topic is configured with `(N)PMSGDLV = ALLAVAIL`. An attempt to use any other values results in a `MQRC_PUBLICATION_FAILURE` error. If the delivery delay processor fails while it is putting the message to the target queue, the result is as described in the "Failure to put to target queue" section.

### **Report messages**

All report options are supported and actioned by the delivery processor, other than the following options that are ignored, but passed through on the message when it is sent to the target queue:

- `MQRO_COA*`
- `MQRO_COD*`
- `MQRO_PAN/MQRO_NAN`
- `MQRO_ACTIVITY`

*Cloned and shared subscriptions:*

In IBM MQ Version 8.0 or later, there are two methods for giving multiple consumers access to the same subscription. These two methods are by using cloned subscriptions, or by using shared subscriptions.

### **Cloned subscriptions**

Cloned subscription is an IBM MQ extension. Cloned subscriptions enable multiple consumers in different Java virtual machines (JVMs) concurrent access to the subscription. This behavior can be used by setting the `CLONESUPP` property to `Enabled` on a `ConnectionFactory` object. By default `CLONESUPP` is `Disabled`. Cloned subscriptions can be enabled only on durable subscriptions. If `CLONESUPP` is enabled, each subsequent connection that is made by using this `ConnectionFactory` is cloned.

A durable subscription can be considered cloned if one or more consumers are created to receive messages from that subscription, that is, they were created specifying the same subscription name. This can be done only if the connection under which the consumers were created has `CLONESUPP` set to `Enabled` on the `MQConnectionFactory`. When a message is published on the subscription's topic, a copy of that message is sent to the subscription. The message is available to any of the consumers, but only one receives it.

**Note:** Enabling cloned subscriptions extends the JMS specification.

## Shared subscriptions

The JMS 2.0 specification introduces shared subscriptions, which allows messages from a topic subscription to be shared among multiple consumers. Each message from the subscription is delivered to only one of the consumers on that subscription. Shared subscriptions are enabled by the relevant call to the JMS 2.0 API.

The APIs can be called in either of the following ways:

- From a Java SE application (or Java EE Client Container).
- From a servlet or the implementation of an MDB.

The JMS 2.0 specification does not define any standard way of driving an MDB from a sharedSubscription, so IBM MQ Version 8.0 or later provides the sharedSubscription activation specification property for this purpose. For more information about this property, see “Configuring the resource adapter for inbound communication” on page 1106 and “Examples of how to define the sharedSubscription property” on page 1120.

If a shared subscription is enabled, then it cannot be unshared.

Shared subscriptions can be created as either durable or non-durable subscriptions. There is no requirement to separately create any objects on the queue manager side beyond the normal JMS configuration, any objects that are required are created dynamically.

## Deciding between shared or cloned subscriptions

When you are determining whether to use shared or cloned subscriptions, consider the benefits of both. Where possible use shared subscriptions as it is specification defined behavior, rather than an IBM MQ specific extension.

The following table contains some of the points to consider when you are deciding between shared and cloned subscriptions:

*Table 146. Considerations when you are choosing between shared subscriptions and cloned subscriptions*

Shared Subscriptions	Cloned Subscriptions
Shared subscriptions are a standard part of the JMS 2.0 specification.	Cloned subscriptions are an IBM MQ specific extension.
Shared subscriptions are created by using explicit API method calls.	Cloned subscriptions are controlled administratively at the ConnectionFactory level.
Shared subscriptions can be durable or nondurable.	Cloned subscriptions can only be durable.
Shared subscriptions are explicitly created on an individual subscription basis.	Cloned subscriptions are used for any durable subscription under a connection for which the function is enabled.
If a subscription is created as shared it cannot later be changed to unshared, or vice versa.	A subscription can be changed from cloned to uncloned each time it is reopened if the <b>CLONESUPP</b> property of the owning connection has changed.

### ► V9.0.0.1

## Attempts to alter the sharability of an existing subscription

If a subscription is created as shared it cannot later be changed to unshared, or vice versa.

From IBM MQ Version 9.0.0, Fix Pack 1, the IBM MQ queue manager has been updated so that if a JMS 2.0 application creates a durable unshared subscription, and another non JMS 2.0 application then tries to resume the subscription, the following reason code is returned:



2432 (MQRC\_SUB\_ALREADY\_EXISTS)

**Related reference:**

“Examples of how to define the sharedSubscription property” on page 1120

You can define the sharedSubscription property of an activation specification within a WebSphere Application Server Liberty server.xml file. Alternatively, you can define the property within a message driven bean (MDB) using annotations.

**Related information:**

Subscribers and subscriptions

Subscription durability

Using JMS 2.0 shared subscriptions

CLONESUPP

*SupportMQExtensions property:*

The JMS 2.0 specification introduces changes to the way certain behaviors work. IBM MQ Version 8.0 and later includes the property com.ibm.mq.jms.SupportMQExtensions, which can be set to *TRUE*, to revert these changed behaviors back to previous implementations.

Three areas of functionality are reverted by setting SupportMQExtensions to *True*:

**Message priority**

Messages can be assigned a priority, 0 - 9. Before JMS 2.0, messages could also use the value -1, indicating that a queue's default priority is used. JMS 2.0 does not allow a message priority of -1 to be set. Turning on SupportMQExtensions allows the value of -1 to be used.

**Client id**

The JMS 2.0 specification requires that non-null client ids are checked for uniqueness when they make a connection. Turning on SupportMQExtensions, means that this requirement is disregarded, and that a client ID can be reused.

**NoLocal**

The JMS 2.0 specification requires that when this constant is turned on, a consumer cannot receive messages that are published by the same client ID. Before JMS 2.0, this attribute was set on a subscriber to prevent it receiving messages that are published by its own connection. Turning on SupportMQExtensions reverts this behavior to its previous implementation.

The property com.ibm.mq.jms.SupportMQExtensions is a boolean property contained in com.ibm.mqjms.jar. This property can be set as follows:

```
java -Dcom.ibm.mq.jms.SupportMQExtensions=true
```

This property can be set either as a standard JVM System property on the **java** command or contained within the IBM MQ classes for JMS configuration file.

**Related concepts:**

“The IBM MQ classes for JMS configuration file” on page 837

An IBM MQ classes for JMS configuration file specifies properties that are used to configure IBM MQ classes for JMS.

**Related reference:**

“Properties used to configure JMS client behavior” on page 843

Use these properties to configure the behavior of the JMS client.

**IBM MQ classes for JMS Application Server Facilities**

This topic describes how IBM MQ classes for JMS implements the `ConnectionConsumer` class and advanced functionality in the `Session` class. It also summarizes the function of a server session pool.

**Important:** This information is for reference only. An application must not be written to use this interface: it is used within the IBM MQ resource adapter to connect to Java EE servers. For practical connection information, see “Using the IBM MQ resource adapter” on page 1092.

IBM MQ classes for JMS supports the Application Server Facilities (ASF) that are specified in the *Java Message Service Specification* (see Oracle Technology Network for Java Developers ). This specification identifies three roles within this programming model:

- **The JMS provider** supplies `ConnectionConsumer` and advanced `Session` functionality.
- **The application server** supplies `ServerSessionPool` and `ServerSession` functionality.
- **The client application** uses the functionality that the JMS provider and application server supply.

The information in this topic does not apply if an application uses a real-time connection to a broker.

**The JMS `ConnectionConsumer`:**

The `ConnectionConsumer` interface provides a high-performance method to deliver messages concurrently to a pool of threads.

The JMS specification enables an application server to integrate closely with a JMS implementation by using the `ConnectionConsumer` interface. This feature provides concurrent processing of messages. Typically, an application server creates a pool of threads, and the JMS implementation makes messages available to these threads. A JMS-aware application server (such as WebSphere Application Server) can use this feature to provide high-level messaging functionality, such as message driven beans.

Normal applications do not use the `ConnectionConsumer`, but expert JMS clients might use it. For such clients, the `ConnectionConsumer` provides a high-performance method to deliver messages concurrently to a pool of threads. When a message arrives on a queue or a topic, JMS selects a thread from the pool and delivers a batch of messages to it. To do this, JMS runs an associated `MessageListener`'s `onMessage()` method.

You can achieve the same effect by constructing multiple `Session` and `MessageConsumer` objects, each with a registered `MessageListener`. However, the `ConnectionConsumer` provides better performance, less use of resources, and greater flexibility. In particular, fewer `Session` objects are required.

## Planning an application with ASF:

This section tells you how to plan an application including:

- “General principles for point-to-point messaging using ASF”
- “General principles for publish/subscribe messaging using ASF” on page 1004
- “Removing messages from the queue in ASF” on page 1005
- Handling poison messages in ASF. See “Handling poison messages in IBM MQ classes for JMS” on page 956.

*General principles for point-to-point messaging using ASF:*

Use this topic for general information about point-to-point messaging using ASF.

When an application creates a `ConnectionConsumer` from a `QueueConnection` object, it specifies a JMS queue object and a selector string. The `ConnectionConsumer` then begins to provide messages to sessions in the associated `ServerSessionPool`. Messages arrive on the queue, and if they match the selector, they are delivered to sessions in the associated `ServerSessionPool`.

In IBM MQ terms, the queue object refers to either a `QLOCAL` or a `QALIAS` on the local queue manager. If it is a `QALIAS`, that `QALIAS` must refer to a `QLOCAL`. The fully resolved IBM MQ `QLOCAL` is known as the *underlying QLOCAL*. A `ConnectionConsumer` is said to be *active* if it is not closed and its parent `QueueConnection` is started.

It is possible for multiple `ConnectionConsumers`, each with different selectors, to run against the same underlying `QLOCAL`. To maintain performance, unwanted messages must not accumulate on the queue. Unwanted messages are those for which no active `ConnectionConsumer` has a matching selector. You can set the `QueueConnectionFactory` so that these unwanted messages are removed from the queue (for details, see “Removing messages from the queue in ASF” on page 1005 ). You can set this behavior in one of two ways:

- Use the JMS administration tool to set the `QueueConnectionFactory` to `MRET(NO)`.
- In your program, use:

```
MQQueueConnectionFactory.setMessageRetention(WMQConstants.WMQ_MRET_NO)
```

If you do not change this setting, the default is to retain such unwanted messages on the queue.

When you set up the IBM MQ queue manager, consider the following points:

- The underlying `QLOCAL` must be enabled for shared input. To do this, use the following MQSC command:  

```
ALTER QLOCAL(your.qlocal.name) SHARE GET(ENABLED)
```
- Your queue manager must have an enabled dead-letter queue. If a `ConnectionConsumer` experiences a problem when it puts a message on the dead-letter queue, message delivery from the underlying `QLOCAL` stops. To define a dead-letter queue, use:  

```
ALTER QMGR DEADQ(your.dead.letter.queue.name)
```
- The user that runs the `ConnectionConsumer` must have authority to perform `MQOPEN` with `MQOO_SAVE_ALL_CONTEXT` and `MQOO_PASS_ALL_CONTEXT`. For details, see the IBM MQ documentation for your specific platform.
- If unwanted messages are left on the queue, they degrade the system performance. Therefore, plan your message selectors so that between them, the `ConnectionConsumers` will remove all messages from the queue.

For details about MQSC commands, see MQSC commands.

*General principles for publish/subscribe messaging using ASF:*

ConnectionConsumers receive messages for a specified Topic. A ConnectionConsumer can be durable or non-durable. You must specify which queue or queues the ConnectionConsumer uses.

When an application creates a ConnectionConsumer from a TopicConnection object, it specifies a Topic object and a selector string. The ConnectionConsumer then begins to receive messages that match the selector on that Topic, including any retained publications for the topic subscribed to.

Alternatively, an application can create a durable ConnectionConsumer that is associated with a specific name. This ConnectionConsumer receives messages that have been published on the Topic since the durable ConnectionConsumer was last active. It receives all such messages that match the selector on the Topic. However, if the ConnectionConsumer is using read-ahead, it can lose nonpersistent messages that are in the client buffer when it closes.

If IBM MQ classes for JMS is in IBM MQ messaging provider migration mode, a separate queue is used for non-durable ConnectionConsumer subscriptions. The CCSUB configurable option on the TopicConnectionFactory specifies the queue to use. Normally, the CCSUB specifies a single queue for use by all ConnectionConsumers that use the same TopicConnectionFactory. However, it is possible to make each ConnectionConsumer generate a temporary queue by specifying a queue name prefix followed by an asterisk (\*).

If IBM MQ classes for JMS is in IBM MQ messaging provider migration mode, the CCDSUB property of the Topic specifies the queue to use for durable subscriptions. Again, this can be a queue that already exists or a queue name prefix followed by an asterisk (\*). If you specify a queue that already exists, all durable ConnectionConsumers that subscribe to the Topic use this queue. If you specify a queue name prefix followed by an asterisk (\*), a queue is generated the first time that a durable ConnectionConsumer is created with a particular name. This queue is reused later when a durable ConnectionConsumer is created with the same name.

When you set up the IBM MQ queue manager, consider the following points:

- Your queue manager must have an enabled dead-letter queue. If a ConnectionConsumer experiences a problem when it puts a message on the dead-letter queue, message delivery from the underlying QLOCAL stops. To define a dead-letter queue, use:  
`ALTER QMGR DEADQ( your.dead.letter.queue.name )`
- The user that runs the ConnectionConsumer must have authority to perform MQOPEN with MQOO\_SAVE\_ALL\_CONTEXT and MQOO\_PASS\_ALL\_CONTEXT. For details, see the IBM MQ documentation for your platform.
- You can optimize performance for an individual ConnectionConsumer by creating a separate, dedicated, queue for it. This is at the cost of extra resource usage.

*Removing messages from the queue in ASF:*

When an application uses `ConnectionConsumers`, JMS might need to remove messages from the queue in a number of situations.

These situations are as follows:

**Badly formatted message**

A message might arrive that JMS cannot parse.

**Poison message**

A message might reach the backout threshold, but the `ConnectionConsumer` fails to requeue it on the backout queue.

**No interested `ConnectionConsumer`**

For point-to-point messaging, when the `QueueConnectionFactory` is set so that it does not retain unwanted messages, a message arrives that is unwanted by any of the `ConnectionConsumers`.

In these situations, the `ConnectionConsumer` attempts to remove the message from the queue. The disposition options in the report field of the message's MQMD set the exact behavior. These options are:

**MQRO\_DEAD\_LETTER\_Q**

The message is requeued to the queue manager's dead-letter queue. This is the default.

**MQRO\_DISCARD\_MSG**

The message is discarded.

The `ConnectionConsumer` also generates a report message, and this also depends on the report field of the message's MQMD. This message is sent to the message's `ReplyToQ` on the `ReplyToQmgr`. If there is an error while the report message is being sent, the message is sent to the dead-letter queue instead. The exception report options in the report field of the message's MQMD set details of the report message. These options are:

**MQRO\_EXCEPTION**

A report message is generated that contains the MQMD of the original message. It does not contain any message body data.

**MQRO\_EXCEPTION\_WITH\_DATA**

A report message is generated that contains the MQMD, any MQ headers, and 100 bytes of body data.

**MQRO\_EXCEPTION\_WITH\_FULL\_DATA**

A report message is generated that contains all data from the original message.

**default**

No report message is generated.

When report messages are generated, the following options are honored:

- `MQRO_NEW_MSG_ID`
- `MQRO_PASS_MSG_ID`
- `MQRO_COPY_MSG_ID_TO_CORREL_ID`
- `MQRO_PASS_CORREL_ID`

If a poison message cannot be requeued, perhaps because the dead-letter queue is full or authorization is wrongly specified, what happens depends on the persistence of the message. If the message is nonpersistent, the message is discarded and no report message is generated. If the message is persistent, delivery of messages to all connection consumers listening on that destination stops. Such connection consumers must be closed and the problem resolved before they can be re-created and message delivery restarted.

It is important to define a dead-letter queue, and to check it regularly to ensure that no problems occur. Particularly, ensure that the dead-letter queue does not reach its maximum depth, and that its maximum message size is large enough for all messages.

When a message is requeued to the dead-letter queue, it is preceded by an IBM MQ dead-letter header (MQDLH). See MQDLH - Dead-letter header for details about the format of the MQDLH. You can identify messages that a ConnectionConsumer has placed on the dead-letter queue, or report messages that a ConnectionConsumer has generated, by the following fields:

- PutApplType is MQAT\_JAVA (0x1C)
- PutApplName is " MQ JMS ConnectionConsumer "

These fields are in the MQDLH of messages on the dead-letter queue, and the MQMD of report messages. The feedback field of the MQMD, and the Reason field of the MQDLH, contain a code describing the error. For details about these codes, see "Reason and feedback codes in ASF" on page 1007. Other fields are as described in MQDLH - Dead-letter header.

#### *Handling poison messages in ASF:*

Within the Application Server Facilities, poison message handling is handled slightly differently to elsewhere in IBM MQ classes for JMS.

For information about poison message handling in IBM MQ classes for JMS, see "Handling poison messages in IBM MQ classes for JMS" on page 956.

When you use Application Server Facilities (ASF), the ConnectionConsumer, rather than the MessageConsumer, processes poison messages. The ConnectionConsumer requeues messages according to the BackoutThreshold and BackoutRequeueQName properties of the queue.

When an application uses ConnectionConsumers, the circumstances in which a message is backed out depend on the session that the application server provides:

- When the session is non-transacted, with AUTO\_ACKNOWLEDGE or DUPS\_OK\_ACKNOWLEDGE, a message is backed out only after a system error, or if the application terminates unexpectedly.
- When the session is non-transacted with CLIENT\_ACKNOWLEDGE, unacknowledged messages can be backed out by the application server calling Session.recover().

Typically, the client implementation of MessageListener or the application server calls Message.acknowledge(). Message.acknowledge() acknowledges all messages delivered on the session so far.

- When the session is transacted, unacknowledged messages can be backed out by the application server calling Session.rollback().
- If the application server supplies an XASession, messages are committed or backed out depending on a distributed transaction. The application server takes responsibility for completing the transaction.

The embedded JMS provider in WebSphere Application Server, Version 5.0 and Version 5.1 handles poison messages in a different way to that just described for IBM MQ classes for JMS. For information about how the embedded JMS provider handles poison messages, see the relevant WebSphere Application Server product documentation.

## Related concepts:

“Handling poison messages in IBM MQ classes for JMS” on page 956

A poison message is one which cannot be processed by a receiving MDB application. If a poison message is encountered, the JMS MessageConsumer and ConnectionConsumer objects can requeue it according to two queue properties, BOQUEUE, and BOTHRESH.

## Error handling:

This section covers various aspects of error handling, including “Recovering from error conditions in the ASF” and “Reason and feedback codes in ASF.”

### *Recovering from error conditions in the ASF:*

If a ConnectionConsumer experiences a serious error, message delivery to all ConnectionConsumers with an interest in the same QLOCAL stops. When this occurs, any ExceptionListener that is registered with the affected Connection is notified. There are two ways in which an application can recover from these error conditions.

Typically, a serious error of this nature occurs if the ConnectionConsumer cannot requeue a message to the dead-letter queue, or it experiences an error when reading messages from the QLOCAL.

Because any ExceptionListener that is registered with the affected Connection is notified, you can use them to identify the cause of the problem. In some cases, the system administrator must intervene to resolve the problem.

Use one of the following techniques to recover from these error conditions:

- Call `close()` on all affected ConnectionConsumers. The application can create new ConnectionConsumers only after all affected ConnectionConsumers are closed and any system problems are resolved.
- Call `stop()` on all affected Connections. After all Connections are stopped and any system problems are resolved, the application can start `()` its Connections successfully.

### *Reason and feedback codes in ASF:*

Use reason and feedback codes to determine the cause of an error. Common reason codes generated by the ConnectionConsumer are given here.

To determine the cause of an error, use the following information:

- The feedback code in any report messages
- The reason code in the MQDLH of any messages in the dead-letter queue

ConnectionConsumers generate the following reason codes.

#### **MQRC\_BACKOUT\_THRESHOLD\_REACHED (0x93A; 2362)**

**Cause** The message has reached the Backout Threshold defined on the QLOCAL, but no Backout Queue is defined.

On platforms where you cannot define the Backout Queue, the message has reached the JMS-defined backout threshold of 20.

#### **Action**

If this is not wanted, define the Backout Queue for the relevant QLOCAL. Also look for the cause of the multiple backouts.

#### **MQRC\_MSG\_NOT\_MATCHED (0x93B; 2363)**

**Cause** In point-to-point messaging, there is a message that does not match any of the selectors

for the ConnectionConsumers monitoring the queue. To maintain performance, the message is requeued to the dead-letter queue.

**Action**

To avoid this situation, ensure that ConnectionConsumers using the queue provide a set of selectors that deal with all messages, or set the QueueConnectionFactory to retain messages.

Alternatively, investigate the source of the message.

**MQRC\_JMS\_FORMAT\_ERROR (0x93C; 2364)**

**Cause** JMS cannot interpret the message on the queue.

**Action**

Investigate the origin of the message. JMS normally delivers messages of an unexpected format as a BytesMessage or TextMessage. Occasionally, this fails if the message is very badly formatted.

Other codes that appear in these fields are caused by a failed attempt to requeue the message to a Backout Queue. In this situation, the code describes the reason that the requeue failed. To diagnose the cause of these errors, refer to API reason codes.

If the report message cannot be put on the ReplyToQ, it is put on the dead-letter queue. In this situation, the feedback field of the MQMD is completed as described in this topic. The reason field in the MQDLH explains why the report message could not be placed on the ReplyToQ.

**The function of a server session pool in AFS:**

This topic summarizes the function of a server session pool.

Figure 145 on page 1009 summarizes the principles of ServerSessionPool and ServerSession functionality.



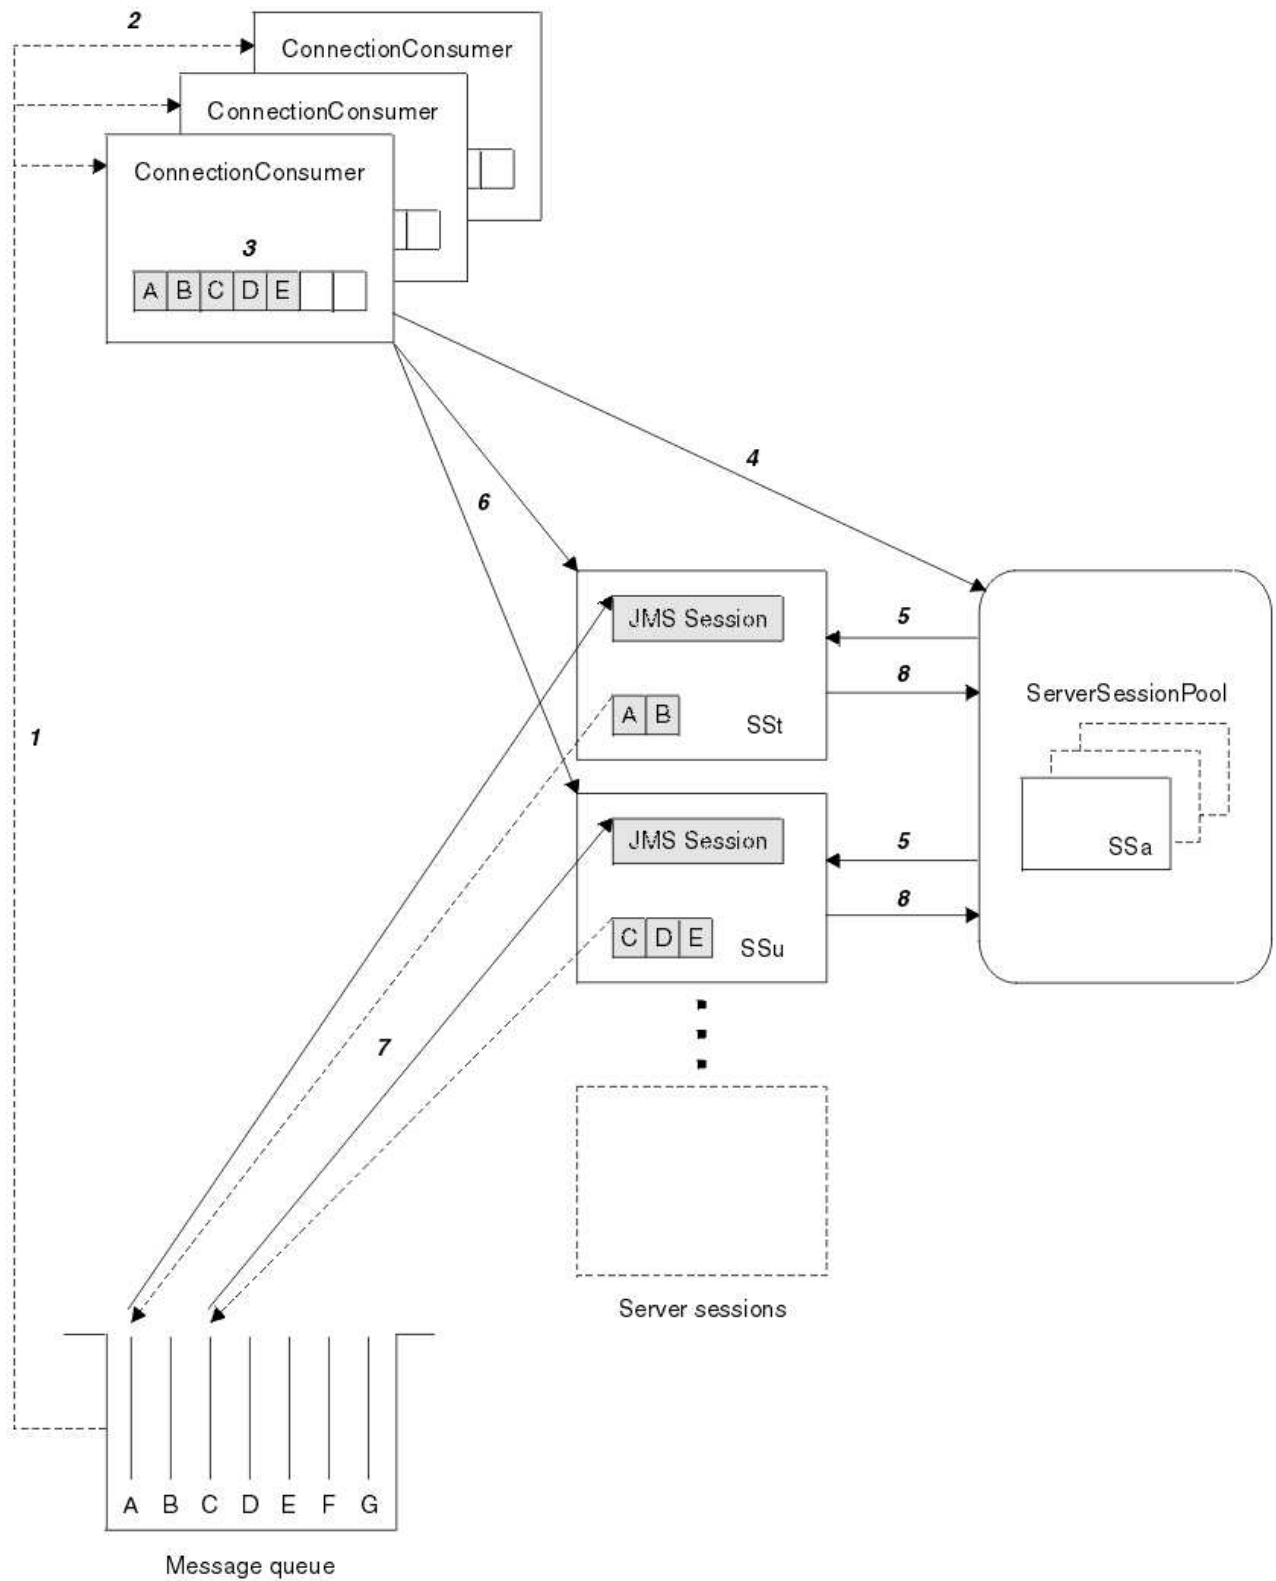


Figure 145. ServerSessionPool and ServerSession functionality

1. The ConnectionConsumers get message references from the queue.
2. Each ConnectionConsumer selects specific message references.
3. The ConnectionConsumer buffer holds the selected message references.

4. The `ConnectionConsumer` requests one or more `ServerSessions` from the `ServerSessionPool`.
5. `ServerSessions` are allocated from the `ServerSessionPool`.
6. The `ConnectionConsumer` assigns message references to the `ServerSessions` and starts the `ServerSession` threads running.
7. Each `ServerSession` retrieves its referenced messages from the queue. It passes them to the `onMessage` method from the `MessageListener` that is associated with the JMS Session.
8. After it completes its processing, the `ServerSession` is returned to the pool.

An application server normally supplies `ServerSessionPool` and `ServerSession` functionality.

## Using IBM MQ classes for JMS in a CICS OSGi JVM server

z/OS

IBM MQ Version 8.0 added support for using the IBM MQ classes for JMS in certain versions of the CICS Open Services Gateway initiative (OSGi) Java Virtual Machine (JVM) server.

**Attention:** Check the system requirements for the CICS system that your enterprise uses. See Detailed system requirements for CICS Transaction Server for detailed information.

This set of topics describe how to set up the IBM MQ classes for JMS in a JVM server environment, and the API restrictions that apply when using the classic (JMS 1.1) and simplified (JMS 2.0) interfaces. See “JMS API restrictions” on page 1013 for a list of the API-specific information.

**Note:** Similar restrictions apply to the legacy (JMS 1.0.2) domain-specific interfaces, but they are not specifically described here.

### General restrictions

The following restrictions apply when using IBM MQ classes for JMS in a CICS OSGi JVM server:

- Client mode connections are not supported.
- Connections are only supported to IBM WebSphere MQ Version 7.1 or IBM MQ Version 8.0 or later queue managers. The **PROVIDERVERSION** attribute on the connection factory must be either unspecified, or a value greater than, or equal to, seven.
- Use of any of the XA connection factories, for example `com.ibm.mq.jms.MQXAConnectionFactory`, is not supported.
- Use of IBM MQ classes for JMS in a CICS OSGi JVM server is only supported in CICS Version 5.2 or later. If you are using CICS Version 5.2, you must apply APAR PI32151.

V 9.0.1

Before IBM MQ Version 9.0.1, using the IBM MQ classes for JMS in a Liberty JVM server environment is not supported.

## Related information:

Configuring the JMS **PROVIDERVERSION** property

## Setting up the JVM server environment:

How you set up your system to run IBM MQ classes for JMS in a CICS OSGi JVM server.

### Procedure

1. Set up the CICS adapter. See Setting up the CICS- IBM MQ adapter.
2. Add the IBM MQ classes for JMS to the JVM server as an OSGi middleware bundle. Do this, by including the following lines in the JVM profile of the JVM server where `MQ_ROOT` is the `java/lib/` directory of the IBM MQ for z/OS Unix System Services installation, for example, `/usr/lpp/V9R0M0/java/lib`.

```
OSGI_BUNDLES=MQ_INSTALLATION_PATH/OSGiCurrent/com.ibm.mq.osgi.allclientprereqs.jar, \
MQ_INSTALLATION_PATH/OSGiCurrent/com.ibm.mq.osgi.allclient.jar
```

3. Add the directory containing the IBM MQ classes for JMS native libraries to the `LIBPATH_SUFFIX` in the JVM profile of the JVM server. For example:

```
LIBPATH_SUFFIX=MQ_INSTALLATION_PATH
```

4. Stop and restart the JVM server.

### What to do next

Set up the CICS Explorer Software Development Kit (SDK) environment.

Do this by following the steps described in Updating the target platform to add

```
MQ_INSTALLATION_PATH/OSGiCurrent/com.ibm.mq.osgi.allclientprereqs.jar
```

and

```
MQ_INSTALLATION_PATH/OSGiCurrent/com.ibm.mq.osgi.allclient.jar
```

to the CICS Explorer SDK.

This allows the IBM MQ classes for JMS to be used by Java applications developed in the CICS Explorer SDK.

## Creating and configuring connection factories and destinations:

There are three possible approaches for instantiating the IBM MQ implementations of connection factories and destinations.

See Creating and configuring connection factories for more information.

These approaches are valid in a Java Virtual Machine (JVM) server environment with the following restrictions:

- If a queue manager, or queue-sharing group, name is specified on a connection factory, the name is ignored.  
The queue manager, or queue-sharing group, that is connected to from CICS is always the one specified in the `MQCONN` resource defined in the CICS region associated with the JVM server.
- If using the Java Naming and Directory Interface (JNDI), the `com.sun.jndi.ldap.LdapCtxFactory` initial context factory is not supported, as it does not work in an OSGi environment.
- If using JNDI, the initial context factory of the JNDI implementation used, and the IBM MQ object factories, must be programmatically registered with OSGi. See “Registering the initial context factory and IBM MQ object factories with OSGi” on page 1012 for further information.

## Registering the initial context factory and IBM MQ object factories with OSGi

In order to use JNDI to locate IBM MQ connection factories and destinations, you need to register both the initial context factory of the JNDI provider, and the object factories provided by IBM MQ as services with the OSGi environment.

Typically, you do this in the start method of a bundle activator. A sample bundle activator that performs this registration for the file system based JNDI repository, provided by fscontext.jar, follows:

```
import javax.naming.spi.*;
import org.osgi.framework.*;
import com.ibm.mq.jms.*;
import com.sun.jndi.fscontext.ReffFSContextFactory;

public class SampleBundleActivator implements BundleActivator
{
 @Override
 public void start(BundleContext ctx) throws Exception
 {
 //Register the initial context factory implementation of the JNDI provider
 ctx.registerService(new String[]{InitialContextFactory.class.getName(),
 ReffFSContextFactory.class.getName()}, new ReffFSContextFactory(), null);

 //Register the MQ object factories.
 final String ofClassName = ObjectFactory.class.getName();
 ctx.registerService(new String[]{ofClassName,
 MQConnectionFactoryFactory.class.getName()}, new MQConnectionFactoryFactory(), null);

 ctx.registerService(new String[]{ofClassName,
 MQQueueConnectionFactoryFactory.class.getName()},
 new MQQueueConnectionFactoryFactory(), null);

 ctx.registerService(new String[]{ofClassName,
 MQTopicConnectionFactoryFactory.class.getName()},
 new MQTopicConnectionFactoryFactory(), null);

 ctx.registerService(new String[]{ofClassName, MQQueueFactory.class.getName()},
 new MQQueueFactory(), null);

 ctx.registerService(new String[]{ofClassName, MQTopicFactory.class.getName()},
 new MQTopicFactory(), null);
 }

 @Override
 public void stop(BundleContext ctx) throws Exception
 {
 }
}
```

### Transactional behavior:

Messages sent and received by the IBM MQ classes for JMS in a JVM server environment are always associated with the CICS unit of work (UOW) that is active on the current thread.

That UOW can only be completed by calling the commit or rollback methods on the `com.ibm.cics.server.Task` object, or by the CICS task ending normally in which case the UOW is implicitly committed.

As a result of this, the values of the **transacted** and **acknowledgeMode** arguments are ignored when calling any of the `Connection.createSession`, or `ConnectionFactory.createContext` methods. Additionally the following methods are not supported. Calling any of the following methods results in an `IllegalStateException` in the session case:

- `javax.jms.Session.commit()`
- `javax.jms.Session.recover()`
- `javax.jms.Session.rollback()`

and an `IllegalStateException` in the JMS context case:

- `javax.jms.JMSContext.commit()`
- `javax.jms.JMSContext.recover()`
- `javax.jms.JMSContext.rollback()`

There is one exception to this behavior. If a session or JMS context is created using one of the following mechanisms:

- `Connection.createSession(false, Session.AUTO_ACKNOWLEDGE)`
- `Connection.createSession(Session.AUTO_ACKNOWLEDGE)`
- `ConnectionFactory.createContext(JMSContext.AUTO_ACKNOWLEDGE)`

then the behavior of that session, or JMS context, is as follows:

- Any messages that are sent are transferred outside of the CICS UOW. That is, they will be available on the target destination immediately, or when the provided delivery delay interval has completed.
- Any non-persistent messages will be received outside of the CICS UOW, unless the `SYNCPOINTALLGETS` property has been specified on the connection factory that created the session or JMS context.
- Persistent messages will always be received inside the CICS UOW.

### Connection authentication and authorization

The JMS specification allows a user name and password to be specified for authentication and authorization when creating a connection or JMS context object.

This is not supported in a JVM server environment. Attempting to create a connection while specifying a user name and password results in a `JMSEException` being thrown. Attempting to create a JMS context, while specifying a user name and password, results in a `JMSRuntimeException` being thrown.

Instead, existing mechanisms for authentication and authorization when connecting to IBM MQ from a CICS environment must be used.

For more information, see *Setting up security on z/OS*. In particular, refer to *User IDs for security checking*, which describes the user IDs that can be used.

### JMS API restrictions:

From a JMS specification perspective, the IBM MQ classes for JMS treat a JVM server as a Java EE compliant application server, that always has a JTA transaction in progress.

For example, you can never call `javax.jms.Session.commit()` in CICS, because the JMS specification states that you can not call it in a JEE EJB, or Web container, while a JTA transaction is in progress.

This results in the following restrictions to the JMS API, in addition to those described in “Transactional behavior” on page 1012.

### Classic API restrictions

- `javax.jms.Connection.createConnectionConsumer(javax.jms.Destination, String, javax.jms.ServerSessionPool, int)` always throws a `JMSEException`.
- `javax.jms.Connection.createDurableConnectionConsumer(javax.jms.Topic, String, String, javax.jms.ServerSessionPool, int)` always throws a `JMSEException`.
- All three variants of `javax.jms.Connection.createSession` always throws a `JMSEException` if the connection already has an existing session active.
- `javax.jms.Connection.createSharedConnectionConsumer(javax.jms.Topic, String, String, javax.jms.ServerSessionPool, int)` always throws a `JMSEException`.

- `javax.jms.Connection.createSharedDurableConnectionConsumer(javax.jms.Topic, String, String, javax.jms.ServerSessionPool, int)` always throws a `JMSEException`.
- `javax.jms.Connection.setClientID()` always throws a `JMSEException`.
- `javax.jms.Connection.setExceptionListener(javax.jms.ExceptionListener)` always throws a `JMSEException`.
- `javax.jms.Connection.stop()` always throws a `JMSEException`.
- `javax.jms.MessageConsumer.setMessageListener(javax.jms.MessageListener)` always throws a `JMSEException`.
- `javax.jms.MessageConsumer.getMessageListener()` always throws a `JMSEException`.
- `javax.jms.MessageProducer.send(javax.jms.Destination, javax.jms.Message, javax.jms.CompletionListener)` always throws a `JMSEException`.
- `javax.jms.MessageProducer.send(javax.jms.Destination, javax.jms.Message, int, int, long, javax.jms.CompletionListener)` always throws a `JMSEException`.
- `javax.jms.MessageProducer.send(javax.jms.Message, int, int, long, javax.jms.CompletionListener)` always throws a `JMSEException`.
- `javax.jms.MessageProducer.send(javax.jms.Message, javax.jms.CompletionListener)` always throws a `JMSEException`.
- `javax.jms.Session.run()` always throws a `JMSRuntimeException`.
- `javax.jms.Session.setMessageListener(javax.jms.MessageListener)` always throws a `JMSEException`.
- `javax.jms.Session.getMessageListener()` always throws a `JMSEException`.

### Simplified API restrictions

- `javax.jms.JMSContext.createContext(int)` always throws a `JMSRuntimeException`.
- `javax.jms.JMSContext.setClientID(String)` always throws a `JMSRuntimeException`.
- `javax.jms.JMSContext.setExceptionListener(javax.jms.ExceptionListener)` always throws a `JMSRuntimeException`.
- `javax.jms.JMSContext.stop()` always throws a `JMSRuntimeException`.
- `javax.jms.JMSProducer.setAsync(javax.jms.CompletionListener)` always throws a `JMSRuntimeException`.
- `javax.jms.JMSConsumer.getMessageListener()` always throws a `JMSRuntimeException`.
- `javax.jms.JMSConsumer.setMessageListener(javax.jms.MessageListener)` always throws a `JMSRuntimeException`.

## Using IBM MQ classes for JMS in a CICS Liberty JVM server



From IBM MQ Version 9.0.1, Java programs running in a CICS Liberty JVM server can use the IBM MQ classes for JMS to access IBM MQ.

You must be using an IBM MQ Version 9.0.1 or later version of the IBM MQ, resource adapter which is obtainable from Fix Central (see “Installing the resource adapter in Liberty” on page 1100).

There are two flavors of Liberty Profile JVMs available in CICS Version 5.3 and later, the types of connection possible to IBM MQ are restricted as follows:

### CICS Liberty Standard

- The IBM MQ resource adapter can connect to any in-service version of IBM MQ in CLIENT mode
- The IBM MQ resource adapter can connect to any in-service version of IBM MQ for z/OS in BINDINGS mode when there is no CICS connection (active CICS MQCONN resource definition) to the same queue manager from the same CICS region.

## CICS Liberty Integrated

- The IBM MQ resource adapter can connect to any in-service version of IBM MQ in CLIENT mode.
- BINDINGS mode connection is not supported.

For more information, see the following topics in the CICS Transaction Server product documentation:

- Accessing IBM MQ from Java programs
- Using IBM MQ classes for JMS in a Liberty JVM server
- Making IBM MQ classes for JMS work in the CICS Liberty IBM MQ environment
- Configuring a Liberty JVM server to support IBM MQ classes for JMS
- Programming with IBM MQ classes for JMS in the CICS IBM MQ environment
- Committing a unit of work in a JMS environment
- CICS abends during the processing of JMS requests

## Using IBM MQ classes for JMS in IMS

IBM MQ Version 8.0.0, Fix Pack 4 added support for using the IBM MQ classes for JMS in IMS versions 13 and later.

**Attention:** Check the system requirements for the IMS system that your enterprise uses. See Hardware and software requirements for IMS for further information.

This set of topics describe how to set up the IBM MQ classes for JMS in an IMS environment, and the API restrictions that apply when using the classic (JMS 1.1) and simplified (JMS 2.0) interfaces. See “JMS API restrictions” on page 1020 for a list of the API-specific information.

**Note:** Similar restrictions apply to the legacy (JMS 1.0.2) domain-specific interfaces, but they are not specifically described here.

## Supported IMS dependent regions

The following dependent region types are supported:

- MPR
- BMP
- IFP
- JMP
- JBP

Unless specifically mentioned in the following topics, IBM MQ classes for JMS behaves the same in all region types.

## Supported Java Virtual Machines

IBM MQ classes for JMS requires Java Platform, Standard Edition 7 (Java SE 7) or later.

## Other restrictions

The following restrictions apply when using IBM MQ classes for JMS in an IMS environment:

- Client mode connections are not supported.
- Connections are only supported to IBM MQ Version 8.0 queue managers using the IBM MQ messaging provider Normal, or Version 8 mode.

The **PROVIDERVERSION** attribute on the connection factory must be either unspecified, or a value greater than, or equal to, seven.

- Use of any of the XA connection factories, for example `com.ibm.mq.jms.MQXAConnectionFactory`, is not supported.

**Related information:**

Defining IBM MQ to IMS

**Setting up the IMS adapter for use with IBM MQ classes for JMS:**

IBM MQ classes for JMS make use of the same IBM MQ-IMS adapter as used by other programming languages. This adapter uses the IMS External Subsystem Attach Facility (ESAF).

**Before you begin**

Before completing the following procedure, you must configure the IMS adapter for the relevant queue managers, and IMS control and dependent regions, as described in Setting up the IMS adapter.

**Attention:** You do not need to perform the step that describes building a dynamic stub, unless you need the dynamic stub for other purposes.

Once you have configured the IMS adapter, carry out the following procedure.

**Procedure**

1. Update the LIBPATH variable in the member of your IMS PROCLIB that is referenced by the ENVIRON parameter in your dependent region JCL (for example, DFSJVMEV) so that it includes the IBM MQ classes for JMS native libraries. That is, the zFS directory that contains `libmqjims.so`. For example, DFSJVMEV might look like the following, where the last line is the directory containing the IBM MQ classes for JMS native libraries:

```
LIBPATH=>
/java/java71_31/J7.1/bin/j9vm:>
/java/java71_31/J7.1/bin:>
/ims13/dbdc/imsjava/classic/lib:>
/ims13/dbdc/imsjava/lib:>
/mqm/V8R0M0/java/lib
```

2. Add the IBM MQ classes for JMS to the class path of the JVM, used by your IMS dependent region, by updating the `java.class.path` option. Do this by following the instructions in DFSJVMMS member of the IMS PROCLIB data set.

For example, you can use the following, where the line in bold indicates the update:

```
-Djava.class.path=/ims13/dbdc/imsjava/imsutm.jar:/ims13/dbdc/imsjava/imsudb.jar:
/mqm/V8R0M0/java/lib/com.ibm.mq.allclient.jarLIBPATH_SUFFIX=MQ_INSTALLATION_PATH
```

**Note:** While there are many different jar files available in the directory containing the IBM MQ classes for JMS, you need only the `com.ibm.mq.allclient.jar` file.

3. Stop and restart the any IMS dependent regions that will be making use of the IBM MQ classes for JMS.

**What to do next**

Create and configure connection factories and destinations.

There are three possible approaches for instantiating the IBM MQ implementations of connection factories and destinations. See “Creating and configuring connection factories and destinations in an IBM MQ classes for JMS application” on page 930for details.

Note that these three approaches are all valid in an IMS environment.



**Related information:**

Setting up the IMS adapter  
Defining IBM MQ to IMS

**Transactional behavior:**

Messages sent and received by the IBM MQ classes for JMS in an IMS environment are always associated with the IMS unit of work (UOW) that is active on the current task.

That UOW can only be completed by calling the commit or rollback methods on an instance of the `com.ibm.ims.dli.tm.Transaction` object, or by the IMS task ending normally in which case the UOW is implicitly committed. If the IMS task ends abnormally, the UOW is rolled back.

As a result of this, the values of the **transacted** and **acknowledgeMode** arguments are ignored when calling any of the `Connection.createSession`, or `ConnectionFactory.createContext` methods. Additionally the following methods are not supported. Calling any of the following methods results in an `IllegalStateException` in the session case:

- `javax.jms.Session.commit()`
- `javax.jms.Session.recover()`
- `javax.jms.Session.rollback()`

and an `IllegalStateException` in the JMS context case:

- `javax.jms.JMSContext.commit()`
- `javax.jms.JMSContext.recover()`
- `javax.jms.JMSContext.rollback()`

There is one exception to this behavior. If a session or JMS context is created using one of the following mechanisms:

- `Connection.createSession(false, Session.AUTO_ACKNOWLEDGE)`
- `Connection.createSession(Session.AUTO_ACKNOWLEDGE)`
- `ConnectionFactory.createContext(JMSContext.AUTO_ACKNOWLEDGE)`

then the behavior of that session, or JMS context, is as follows:

- Any messages sent, are sent outside of the IMS UOW. That is, they will be available on the target destination immediately, or when the provided delivery delay interval has completed.
- Any non-persistent messages will be received outside of the IMS UOW, unless the `SYNCPOINTALLGETS` property has been specified on the connection factory that created the session or JMS context.
- Persistent messages will always be received inside the IMS UOW.

This might be useful if, for example, you want to write an audit message to a queue even if the UOW rolls back.

## Implications of IMS syncpoints:

The IBM MQ classes for JMS build upon the existing IBM MQ adapter support which makes use of ESAF. This means that the documented behavior applies, including all open handles being closed by the IMS adapter when a syncpoint occurs.

See “Syncpoints in IMS applications” on page 59 for more information.

To illustrate this point, consider the following code running in a JMP environment. The second call to `mp.send()` results in a `JMSEException` as the `messageQueue.getUnique(inputMessage)` code results in all open IBM MQ connection and object handles being closed.

Similar behavior is observed if the `getUnique()` call was replaced with `Transaction.commit()`, but not if `Transaction.rollback()` was used.

```
//Create a connection to queue manager MQ21.
MQConnectionFactory cf = new MQConnectionFactory();
cf.setQueueManager("MQ21");

Connection c = cf.createConnection();
Session s = c.createSession();

//Send a message to MQ queue Q1.
Queue q = new MQQueue("Q1");
MessageProducer mp = s.createProducer(q);
TextMessage m = s.createTextMessage("Hello world!");
mp.send(m);

//Get a message from an IMS message queue. This results in a GU call
//which results in all MQ handles being closed.
Application a = ApplicationFactory.createApplication();
MessageQueue messageQueue = a.getMessageQueue();
IOMessage inputMessage = a.getIOMessage(MESSAGE_CLASS_NAME);
messageQueue.getUnique(inputMessage);

//This attempt to send another message will result in a JMSEException containing a
//MQRC_HCONN_ERROR as the connection/handle has been closed.
mp.send(m);
```

The correct code to use in this scenario is as follows. In this case the connection to IBM MQ is closed prior to calling `getUnique()`. The connection and session are then re-created in order to send another message.

```
//Create a connection to queue manager MQ21.
MQConnectionFactory cf = new MQConnectionFactory();
cf.setQueueManager("MQ21");

Connection c = cf.createConnection();
Session s = c.createSession();

//Send a message to MQ queue Q1.
Queue q = new MQQueue("Q1");
MessageProducer mp = s.createProducer(q);
TextMessage m = s.createTextMessage("Hello world!");
mp.send(m);

//Close the connection to MQ, which closes all MQ object handles.
//The send of the message will be committed by the subsequent GU call.
c.close();
c = null;
s = null;
mp = null;

//Get a message from an IMS message queue. This results in a GU call.
Application a = ApplicationFactory.createApplication();
MessageQueue messageQueue = a.getMessageQueue();
IOMessage inputMessage = a.getIOMessage(MESSAGE_CLASS_NAME);
messageQueue.getUnique(inputMessage);

//Re-create the connection to MQ and send another message;
c = cf.createConnection();
```

```
s = c.createSession();
mp = s.createProducer(q);
m = s.createTextMessage("Hello world 2!");
mp.send(m);
```

### Considerations when using the IMS adapter:

You need to be aware of the following restrictions. You can have only one connection handle for each queue manager. There are implications in the interaction with IBM MQ when using both JMS and native code. There are limitations to connection authentication and authorization.

#### One connection handle for each queue manager

Only one connection handle at a time to a specific queue manager is allowed in IMS dependent regions. Any subsequent attempts to connect to the same queue manager reuse the existing handle.

While this behavior should not cause any problems in an application that only uses the IBM MQ classes for JMS, this behavior can cause problems in applications that interact with IBM MQ, when using both the IBM MQ classes for JMS and the MQI in native code written in languages, such as COBOL or C.

#### Implications of interacting with IBM MQ when using both JMS and native code

Problems can occur when interleaving Java code and native code that both use IBM MQ functionality and when the connection to IBM MQ is not closed before leaving either the native or Java code.

For example, in the following pseudo code, a connection handle to a queue manager is originally established in Java code using the IBM MQ classes for JMS. The connection handle is reused in COBOL code and invalidated by a call to MQDISC.

The next time the IBM MQ classes for JMS make use of the connection handle a `JMSException` with a reason code of `MQRC_HCONN_ERROR` results.

```
COBOL code running in message processing region
 Use the Java Native Interface (JNI) to call Java code
 Create MQ connection and session - this creates an MQ connection handle
 Send message to MQ queue
 Store connection and session in static variable
 Return to COBOL code

MQCONN - picks up MQ connection handle established in Java code
MQDISC - invalidates connection handle

Use the Java Native Interface (JNI) to call Java code
 Get session from static variable
 Create a message consumer - fails as connection handle invalidated
```

There are other similar usage patterns which can result in `MQRC_HCONN_ERROR`.

While it is possible to share IBM MQ connection handles between native and Java code (for example, the previous example would work if there had not been an `MQDISC` call) in general, the best practice is to close any connection handles before changing from Java to native code, or the other way round.

#### Connection authentication and authorization

The JMS specification allows a user name and password to be specified for authentication and authorization when creating a connection or JMS context object.

This is not supported in an IMS environment. Attempting to create a connection while specifying a user name and password results in a `JMSException` being thrown. Attempting to create a JMS context, while specifying a user name and password, results in a `JMSRuntimeException` being thrown.

Instead, existing mechanisms for authentication and authorization when connecting to IBM MQ from an IMS environment must be used.

For more information, see *Setting up security on z/OS*. In particular, refer to *User IDs for security checking*, which describes the user IDs that can be used.

**Related information:**

*Setting up security on z/OS*

**JMS API restrictions:**

From a JMS specification perspective, the IBM MQ classes for JMS treat IMS as a Java EE compliant application server, that always has a JTA transaction in progress.

For example, you can never call `javax.jms.Session.commit()` in IMS, because the JMS specification states that you can not call it in a JEE EJB, or Web container, while a JTA transaction is in progress.

This results in the following restrictions to the JMS API, in addition to those described in “Transactional behavior” on page 1017.

**Classic API restrictions**

- `javax.jms.Connection.createConnectionConsumer(javax.jms.Destination, String, javax.jms.ServerSessionPool, int)` always throws a `JMSEException`.
- `javax.jms.Connection.createDurableConnectionConsumer(javax.jms.Topic, String, String, javax.jms.ServerSessionPool, int)` always throws a `JMSEException`.
- All three variants of `javax.jms.Connection.createSession` always throws a `JMSEException` if the connection already has an existing session active.
- `javax.jms.Connection.createSharedConnectionConsumer(javax.jms.Topic, String, String, javax.jms.ServerSessionPool, int)` always throws a `JMSEException`.
- `javax.jms.Connection.createSharedDurableConnectionConsumer(javax.jms.Topic, String, String, javax.jms.ServerSessionPool, int)` always throws a `JMSEException`.
- `javax.jms.Connection.setClientID()` always throws a `JMSEException`.
- `javax.jms.Connection.setExceptionListener(javax.jms.ExceptionListener)` always throws a `JMSEException`.
- `javax.jms.Connection.stop()` always throws a `JMSEException`.
- `javax.jms.MessageConsumer.setMessageListener(javax.jms.MessageListener)` always throws a `JMSEException`.
- `javax.jms.MessageConsumer.getMessageListener()` always throws a `JMSEException`.
- `javax.jms.MessageProducer.send(javax.jms.Destination, javax.jms.Message, javax.jms.CompletionListener)` always throws a `JMSEException`.
- `javax.jms.MessageProducer.send(javax.jms.Destination, javax.jms.Message, int, int, long, javax.jms.CompletionListener)` always throws a `JMSEException`.
- `javax.jms.MessageProducer.send(javax.jms.Message, int, int, long, javax.jms.CompletionListener)` always throws a `JMSEException`.
- `javax.jms.MessageProducer.send(javax.jms.Message, javax.jms.CompletionListener)` always throws a `JMSEException`.
- `javax.jms.Session.run()` always throws a `JMSRuntimeException`.
- `javax.jms.Session.setMessageListener(javax.jms.MessageListener)` always throws a `JMSEException`.
- `javax.jms.Session.getMessageListener()` always throws a `JMSEException`.

### **Simplified API restrictions**

- `javax.jms.JMSContext.createContext(int)` always throws a `JMSRuntimeException`.
- `javax.jms.JMSContext.setClientID(String)` always throws a `JMSRuntimeException`.
- `javax.jms.JMSContext.setExceptionListener(javax.jms.ExceptionListener)` always throws a `JMSRuntimeException`.
- `javax.jms.JMSContext.stop()` always throws a `JMSRuntimeException`.
- `javax.jms.JMSProducer.setAsync(javax.jms.CompletionListener)` always throws a `JMSRuntimeException`.

## **Using IBM MQ classes for Java**

Use IBM MQ in a Java environment. IBM MQ classes for Java allow a Java application to connect to IBM MQ as an IBM MQ client, or connect directly to an IBM MQ queue manager.

**Note:** IBM MQ classes for Java are not supported in IMS.

IBM will make no further enhancements to the IBM MQ classes for Java and they are functionally stabilized at the level shipped in IBM MQ Version 8.0.

IBM MQ classes for Java is one of two alternative APIs that Java applications can use to access IBM MQ resources. The other API is IBM MQ classes for JMS.

From IBM MQ Version 8.0, the IBM MQ classes for Java are built with Java 7.

The Java 7 runtime environment supports running earlier class file versions.

IBM MQ classes for Java encapsulate the Message Queue Interface (MQI), the native IBM MQ API and use a similar object model to the C++ and .NET interfaces to IBM MQ.

Programmable options allow IBM MQ classes for Java to connect to IBM MQ in either of the following ways:

- In client mode as an IBM MQ MQI client by using Transmission Control Protocol/Internet Protocol (TCP/IP)
- In bindings mode, connecting directly to IBM MQ by using the Java Native Interface (JNI)

### **Client mode connection**

An IBM MQ classes for Java application can connect to any supported queue manager by using client mode.

To connect to a queue manager in client mode, an IBM MQ classes for Java application can run on the same system on which the queue manager is running, or on a different system. In each case, IBM MQ classes for Java connects to the queue manager over TCP/IP.

For more information on how to write applications to use client mode connections, see “IBM MQ classes for Java connection modes” on page 1043.

### **Bindings mode connection**

When used in bindings mode, IBM MQ classes for Java uses the Java Native Interface (JNI) to call directly into the existing queue manager API, rather than communicating through a network. In most environments, connecting in bindings mode provides better performance for IBM MQ classes for Java applications than connecting in client mode, by avoiding the cost of TCP/IP communication.

Applications that use the IBM MQ classes for Java to connect in bindings mode must run on the same system as the queue manager to which they are connecting.

The Java Runtime Environment, that is being used to run the IBM MQ classes for Java application, must be configured to load the IBM MQ classes for Java libraries; see “IBM MQ classes for Java libraries” on page 1030 for further information.

For more information on how to write applications to use bindings mode connections, see “IBM MQ classes for Java connection modes” on page 1043.

**Related information:**

IBM MQ Java language interfaces

Tracing IBM MQ classes for Java applications

Java and JMS troubleshooting

Using IBM MQ classes for JMS

### Why should I use IBM MQ classes for Java?

A Java application can use either IBM MQ classes for Java or IBM MQ classes for JMS to access IBM MQ resources.

**Note:** Although existing applications that use the IBM MQ classes for Java continue to be fully supported, new applications should use the IBM MQ classes for JMS. Features that have recently been added to IBM MQ, such as asynchronous consume and automatic reconnection, are not available in the IBM MQ classes for Java, but are available in the IBM MQ classes for JMS. For more information, see “Why should I use IBM MQ classes for JMS?” on page 825.

**Note:** IBM MQ classes for Java are functionally stabilized at the level shipped in IBM MQ Version 8.0. For more information, see Stabilization of IBM MQ classes for Java.



### Prerequisites for IBM MQ classes for Java

To use IBM MQ classes for Java, you need certain other software products.

For information about prerequisites for IBM MQ classes for Java, see the System Requirements for IBM MQ web page.

To develop IBM MQ classes for Java applications, you need a Java Development Kit (JDK). Details of the JDKs supported with your operating system can be found in the System Requirements for IBM MQ information.

To run IBM MQ classes for Java applications, you need the following software components:

- An IBM MQ queue manager, for applications that connect to a queue manager
- A Java Runtime Environment (JRE), for each system on which you run applications. A suitable JRE is supplied with IBM MQ.
-  For IBM i, QShell, which is option 30 of the operating system
-  For z/OS, UNIX and Linux System Services (USS)

If you require TLS connections to use cryptographic modules that have been FIPS 140-2 certified, you need the IBM Java JSSE FIPS provider (IBMJSSEFIPS). Every IBM JDK and JRE at Version 1.4.2 or later contains IBMJSSEFIPS.

You can use Internet Protocol Version 6 (IPv6) addresses in your IBM MQ classes for Java applications if IPv6 is supported by your Java virtual machine (JVM) and the TCP/IP implementation on your operating system.

## Running IBM MQ classes for Java applications within Java EE

There are certain restrictions and design considerations that must be taken into account before using IBM MQ classes for Java in Java EE.

IBM MQ classes for Java has restrictions when used within a Java Platform, Enterprise Edition (Java EE) environment. There are also additional considerations that must be taken into account when designing, implementing, and managing an IBM MQ classes for Java application that runs inside a Java EE environment. These restrictions and considerations are outlined in the following sections.

### JTA transactions restrictions

The only supported transaction manager for applications using IBM MQ classes for Java is IBM MQ itself. Although an application under JTA control can make use IBM MQ classes for Java, any work performed through these classes is not controlled by the JTA units of work. They instead form local units of work separate from those managed by the application server through the JTA interfaces. In particular, any rollback of the JTA transaction does not result in a rollback of any sent or received messages. This restriction applies to application or bean managed transactions and to container managed transactions, and all Java EE containers. To perform messaging work directly with IBM MQ inside application server-coordinated transactions, IBM MQ classes for JMS must be used instead.

### Thread creation

IBM MQ classes for Java creates threads internally for various operations. For example, when running in BINDINGS mode to call directly on a local queue manager, the calls are made on a 'worker' thread created internally by IBM MQ classes for Java. Other threads can be created internally, for example to clear unused connections from a connection pool or to remove subscriptions for terminated publish/subscribe applications.

Some Java EE applications (for example those running in EJB and Web containers) must not create new threads. Instead, all work must be performed on the main application threads managed by the application server. When applications use IBM MQ classes for Java, the application server might not be able to distinguish between application code and the IBM MQ classes for Java code, so the threads previously described cause the application to be non-compliant with the container specification. IBM MQ classes for JMS does not break these Java EE specifications and so can be used instead.

### Security restrictions

Security policies implemented by an application server might prevent certain operations that are undertaken by the IBM MQ classes for Java API, such as creating and operating new threads of control (as described in the preceding sections).

For example, application servers typically run with Java Security disabled by default, and allow it to be enabled through some application server-specific configuration (some application servers also allow more detailed configuration of the policies used within Java Security). When Java Security is enabled, IBM MQ classes for Java might break the Java Security policy threading rules defined for the application server, and the API might not be able to create all the threads that it needs in order to function. To prevent problems with thread management, the use of IBM MQ classes for Java is not supported in environments where Java Security is enabled.

### Application isolation considerations

An intended benefit of running applications within a Java EE environment is application isolation. The design and implementation of IBM MQ classes for Java predate the Java EE environment. IBM MQ classes for Java can be used in a manner which does not support the concept of application isolation. Specific examples of considerations in this area include:

- The use of static (JVM process-wide) settings within the MQEnvironment class, such as:

- the user ID and password to be used for connection identification and authentication
- the host name, port, and channel used for client connections
- TLS configuration for secured client connections

Modifying any of the MQEnvironment properties for the benefit of one application also affect other applications using the same properties. When running in a multi-application environment such as Java EE, each application must use its own distinct configuration through the creation of MQQueueManager objects with a specific set of properties, rather than defaulting to the properties configured in the process-wide MQEnvironment class.

- The MQEnvironment class introduces a number of static methods which act globally on all applications using IBM MQ classes for Java within the same JVM process, and there is no way to override this behavior for particular applications. Examples include:
  - configuring TLS properties, such as the location of the keystore
  - configuring client channel exits
  - enabling or disabling diagnostic tracing
  - managing the default connection pool used to optimize the use of connections to queue managers

Invoking such methods affects all applications running in the same Java EE environment.

- Connection pooling is enabled to optimize the process of making multiple connections to the same queue manager. The default connection pool manager is process-wide, and shared by multiple applications. Changes to connection pool configuration, such as replacing the default connection manager for one application using the MQEnvironment.setDefaultConnectionFactory() method therefore affects other applications running in the same Java EE application server.
- TLS is configured for applications using IBM MQ classes for Java using the MQEnvironment class and MQQueueManager object properties. It is not integrated with the managed security configuration of the application server itself. You must ensure that you configure IBM MQ classes for Java appropriately to provide your required level of security, and not use the application server configuration.

## Bindings mode restrictions

IBM MQ and WebSphere Application Server can be installed on the same machine such that the major versions of the queue manager and of the IBM MQ resource adapter (RA) shipped in WebSphere Application Server are different. For instance WebSphere Application Server Version 7.0, which ships an IBM MQ RA level of 7.0.1, can be installed on the same machine as a Version 6 queue manager.

If the queue manager and resource adapter major versions are different, bindings connections cannot be used. Any connections from WebSphere Application Server to the queue manager using the resource adapter must use client type connections. Bindings connections can be used if the versions are the same.



## Character string conversions in IBM MQ classes for Java

From IBM MQ Version 8.0, some of the default behavior regarding character string conversion in the IBM MQ classes for Java has changed.

Before IBM MQ Version 8.0, string conversions in IBM MQ classes for Java was done by calling the `java.nio.charset.Charset.decode(ByteBuffer)` and `Charset.encode(CharBuffer)` methods.

Using either of these methods results in a default replacement ( `REPLACE`) of malformed or untranslatable data.

This behavior can obscure errors in applications, and lead to unexpected characters, for example `?`, in translated data. From IBM MQ Version 8.0, to detect such issues earlier and more effectively, the IBM MQ classes for Java use `CharsetEncoders` and `CharsetDecoders` directly and configure the handling of malformed and untranslatable data explicitly.

From IBM MQ Version 8.0, the default behavior is to `REPORT` such issues by throwing a suitable `MQException` or `JMSEException`.

### Configuring

Translating from UTF-16 (the character representation used in Java) to a native character set, such as UTF-8, is termed encoding, while translating in the opposite direction is termed decoding.

Currently, decoding takes the default behavior for `CharsetDecoders`, reporting errors by throwing an exception.

One setting is used to specify a `java.nio.charset.CodingErrorAction` to control error handling on both encoding and decoding. One other setting is used to control the replacement byte, or bytes, when encoding. The default Java replacement String will be used in decoding operations.

### IBM MQ Classes for Java

```
com.ibm.mq.MQMD
```

Two new fields have been added:

#### `byte[] unMappableReplacement`

The byte sequence that will be written to an encoded string if an input character cannot be translated, and you have specified `REPLACE`.

**Default:** `"?".getBytes()`

The default Java replacement String is used in decoding operations.

#### `java.nio.charset.CodingErrorAction unmappableAction`

Specifies the action to be taken for untranslatable data on encoding and decoding:

**Default:** `CodingErrorAction.REPORT;`

### Setting system defaults

From IBM MQ Version 8.0, the following two Java system properties are available to configure default behavior regarding character string conversion. Note that destinations or messages override these properties.

#### `com.ibm.mq.cfg.jmqi.UnmappableCharacterAction`

Specifies the action to be taken for untranslatable data on encoding and decoding. The value can be `REPORT`, `REPLACE`, or `IGNORE`.

### **com.ibm.mq.cfg.jmqi.UnmappableCharacterReplacement**

Sets or gets the replacement bytes to apply when a character cannot be mapped in an encoding operation. The default Java replacement string is used in decoding operations.

To avoid confusion between Java character and native byte representations, you should specify `com.ibm.mq.cfg.jmqi.UnmappableCharacterReplacement` as a decimal number representing the replacement byte in the native character set.

For example, the decimal value of `?`, as a native byte, is 63 if the native character set is ASCII-based, such as ISO-8859-1, while it is 111 if the native character set is EBCDIC.

#### **Related concepts:**

“Character string conversions in IBM MQ classes for JMS” on page 869

From IBM MQ Version 8.0, some of the default behavior regarding character string conversion in the IBM MQ classes for JMS has changed.



## **Installation and configuration of IBM MQ classes for Java**

This section describes the directories and files that are created when you install IBM MQ classes for Java, and tells you how to configure IBM MQ classes for Java after installation.

### **What is installed for IBM MQ classes for Java:**

The latest version of IBM MQ classes for Java is installed with IBM MQ. You might need to override default installation options to make sure this is done.

For more information about installing IBM MQ see:

-  Installing IBM MQ
-  Installing the IBM MQ for z/OS product



IBM MQ classes for Java are contained in the Java archive (JAR) files, `com.ibm.mq.jar`, and `com.ibm.mq.jmqi.jar`.

Support for standard message headers, such as Programmable Command Format (PCF), is contained in the JAR file `com.ibm.mq.headers.jar`.

Support for Programmable Command Format (PCF) is contained in the JAR file `com.ibm.mq.pcf.jar`.

**Note:** It is not recommended to use the IBM MQ classes for Java within an application server. For information about the restrictions that apply when running in this environment, see “Running IBM MQ classes for Java applications within Java EE” on page 1023. For more information, see Using WebSphere MQ Java Interfaces in J2EE/JEE Environments.

**Important:** Apart from the relocatable JAR files described in this topic, copying the IBM MQ classes for Java JAR files or native libraries to other machines, or to a different location on a machine where the IBM MQ classes for Java have been installed, is not supported. In addition to this, including the `com.ibm.mq.allclient.jar` file, or the IBM MQ classes for Java, within application archives (such as enterprise application archives, or EAR files), is not supported.

  The `JSON4J.jar` file and `com.ibm.msg.client.mqlight` package are not needed by the IBM MQ classes for Java and IBM MQ classes for JMS. From IBM MQ Version 9.0.0, Fix Pack 3 and IBM MQ Version 9.0.5, the following changes are therefore made to the `com.ibm.mq.allclient.jar` file :

- The reference to `JSON4J.jar` file is removed from the class path statement within the manifest file for the `com.ibm.mq.allclient.jar` file.

- The package `com.ibm.msg.client.mqlight` is no longer included inside the `com.ibm.mq.allclient.jar` file.

## Relocatable JAR files

Within an enterprise, the following files can be moved to systems that need to run IBM MQ classes for Java applications:

- `-com.ibm.mq.allclient.jar`
- `-com.ibm.mq.traceControl.jar`
- The Bouncy Castle security provider and CMS support JAR files

The Bouncy Castle security provider and CMS support JAR files are required. For more information, see [Support for non-IBM JREs](#).

The file `com.ibm.mq.allclient.jar` contains the IBM MQ classes for JMS, the IBM MQ classes for Java, and the PCF and Headers Classes. If you move this file to a new location, make sure that you take steps to keep this new location maintained with new IBM MQ Fix Packs. Also, make sure that the use of this file is made known to IBM Support if you are getting an interim fix.

To determine the version of the `com.ibm.mq.allclient.jar` file, use the command:

```
java -jar com.ibm.mq.allclient.jar
```

The following example shows some sample output from this command:

```
C:\Program Files\IBM\MQ_1\java\lib>java -jar com.ibm.mq.allclient.jar
Name: Java Message Service Client
Version: 9.0.0.0
Level: p000-L140428.1
Build Type: Production
Location: file:/C:/Program Files/IBM/MQ_1/java/lib/com.ibm.mq.allclient.jar

Name: WebSphere MQ classes for Java Message Service
Version: 9.0.0.0
Level: p000-L140428.1
Build Type: Production
Location: file:/C:/Program Files/IBM/MQ_1/java/lib/com.ibm.mq.allclient.jar

Name: WebSphere MQ JMS Provider
Version: 9.0.0.0
Level: p000-L140428.1 mqjbd=p000-L140428.1
Build Type: Production
Location: file:/C:/Program Files/IBM/MQ_1/java/lib/com.ibm.mq.allclient.jar

Name: Common Services for Java Platform, Standard Edition
Version: 9.0.0.0
Level: p000-L140428.1
Build Type: Production
Location: file:/C:/Program Files/IBM/MQ_1/java/lib/com.ibm.mq.allclient.jar
```

The `com.ibm.mq.traceControl.jar` file is used to dynamically control trace for IBM MQ classes for JMS applications. For more information, see [Controlling trace in a running process by using IBM MQ classes for Java and IBM MQ classes for JMS](#).








Installation directories for IBM MQ classes for Java:

IBM MQ classes for Java files and samples are installed in different locations according to platform. The location of the Java Runtime Environment (JRE) that is installed with IBM MQ also varies according to the platform.

### Installation directories for IBM MQ classes for Java files

Table 147 shows where the IBM MQ classes for Java files are installed.

Table 147. IBM MQ classes for Java installation directories








Platform	Directory
 AIX	<code>MQ_INSTALLATION_PATH/java/lib</code>
   Solaris	<code>MQ_INSTALLATION_PATH/java/lib</code>
 IBM i	<code>/QIBM/ProdData/mqm/java/lib</code>
 Windows	<code>MQ_INSTALLATION_PATH\java\lib</code>
 z/OS	<code>MQ_INSTALLATION_PATH/mqm/V8R0M0/java /lib</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

### Installation directories for samples

Some sample applications, such as the Installation Verification Programs (IVPs), are supplied with IBM MQ. Table 148 shows where the sample applications are installed. The IBM MQ classes for Java samples are in a subdirectory called `wmqjava`. The PCF samples are in a subdirectory called `pcf`.

Table 148. Samples directories

Platform	Directory
 AIX	<code>MQ_INSTALLATION_PATH/samp/wmqjava/</code>
   Solaris	<code>MQ_INSTALLATION_PATH/samp/wmqjava/</code>
 IBM i	<code>/QIBM/ProdData/mqm/java/samples</code>
 Windows	<code>MQ_INSTALLATION_PATH\tools\wmqjava\</code>
 z/OS	<code>MQ_INSTALLATION_PATH/mqm/V8R0M0/java/samples</code>








`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

### Installation directories for JRE

The IBM MQ classes for JMS require a Java 7 (or above) Java Runtime Environment (JRE). A suitable JRE is installed with IBM MQ. Table 149 on page 1029 shows where this JRE is installed. To run Java programs such as the provided samples, using this JRE, either explicitly invoke `JRE_LOCATION/bin/java` or add `JRE_LOCATION/bin` to the `PATH` environment (or equivalent) for your platform, where `JRE_LOCATION` is

the directory given in Table 149.

Table 149. JRE directories

Platform	Directory
 AIX	<code>MQ_INSTALLATION_PATH/java/jre</code>
   HP-UX, Linux, and Solaris	<code>MQ_INSTALLATION_PATH/java/jre</code>
 IBM i	<code>/QIBM/ProdData/mqm/java/jre</code>
 Windows	<code>MQ_INSTALLATION_PATH\java\jre</code>
 z/OS	<code>MQ_INSTALLATION_PATH/mqm/V8R0M0/java/jre</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Environment variables relevant to IBM MQ classes for Java:








Some environment variables must be set. Specimen values for classpath are given.

For IBM MQ classes for Java, applications to run, their class path must include the appropriate IBM MQ classes for Java directory. To run the sample applications, the class path must also include the appropriate samples directories. This information can be provided in the Java invocation command or in the CLASSPATH environment variable.

**Important:** Setting the Java option `-Xbootclasspath`, to include the IBM MQ classes for Java, is not supported.

Table 150 shows the appropriate CLASSPATH setting to use on each platform to run IBM MQ classes for Java applications, including the sample applications.

Table 150. CLASSPATH setting to run IBM MQ classes for Java applications, including the IBM MQ classes for Java sample applications

Platform	CLASSPATH setting
 AIX	<code>CLASSPATH= MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.jar: MQ_INSTALLATION_PATH/samp/wmqjava/samples:</code>
   HP-UX, Linux, and Solaris	<code>CLASSPATH= MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.jar: MQ_INSTALLATION_PATH/samp/wmqjava/samples:</code>
 IBM i	<code>CLASSPATH=/QIBM/ProdData/mqm/java/lib/com.ibm.mq.jar: /QIBM/ProdData/mqm/java/samples/wmqjava/samples:</code>
 Windows	<code>CLASSPATH= MQ_INSTALLATION_PATH\Java\lib\com.ibm.mq.jar; MQ_INSTALLATION_PATH\tools\wmqjava\samples;</code>
 z/OS	<code>CLASSPATH= MQ_INSTALLATION_PATH/mqm/V8R0M0/java/lib/com.ibm.mq.jar: MQ_INSTALLATION_PATH/mqm/V8R0M0/java/samples/wmqjava: MQ_INSTALLATION_PATH/mqm/V8R0M0/java/samples/pcf</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

If you compile using the `-Xlint` option, you might see a message warning you that `com.ibm.mq.ese.jar` is not present. You can ignore the warning. This file is only present if you have installed Advanced Message Security.

The scripts provided with IBM MQ classes for JMS use the following environment variables:

#### **MQ\_JAVA\_DATA\_PATH**

This environment variable specifies the directory for log and trace output.

#### **MQ\_JAVA\_INSTALL\_PATH**

This environment variable specifies the directory where IBM MQ classes for Java are installed, as shown in IBM MQ classes for Java installation directories.

#### **MQ\_JAVA\_LIB\_PATH**

This environment variable specifies the directory where the IBM MQ classes for Java libraries are stored, as shown in The location of the IBM MQ classes for Java libraries for each platform. Some scripts supplied with IBM MQ classes for Java, such as `IVTRun`, use this environment variable.

**Windows** On Windows, all the environment variables are set automatically during installation.

**UNIX** On UNIX, you can use the script `setjmsenv` (if you are using a 32-bit JVM) or `setjmsenv64` (if you are using a 64-bit JVM) to set the environment variables.

**UNIX** **Linux** On UNIX and Linux, these scripts are in the `MQ_INSTALLATION_PATH/java/bin` directory.

**IBM i** On IBM i, the environment variable `QIBM_MULTI_THREADED` must be set to `Y`. You can then run multithreaded applications in the same way that you run single threaded applications. See [Setting up IBM MQ with Java and JMS](#) for more information.

IBM MQ classes for Java require a Java 7 Java Runtime Environment (JRE). For information about the location of a suitable JRE that is installed with IBM MQ, see “Installation directories for IBM MQ classes for Java” on page 1028.

*IBM MQ classes for Java libraries:*


The location of the IBM MQ classes for Java libraries varies according to platform. Specify this location when you start an application.

To specify the location of the Java Native Interface (JNI) libraries, start your application by using a **java** command with the following format:








```
java -Djava.library.path= library_path application_name
```

where *library\_path* is the path to the IBM MQ classes for Java, which include the JNI libraries. Table 151 on page 1031 shows the location of the IBM MQ classes for Java libraries for each platform. In this table, `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

Table 151. The location of the IBM MQ classes for Java libraries for each platform

Platform	Directory containing the IBM MQ classes for Java libraries
AIX	<code>MQ_INSTALLATION_PATH/java/lib</code> (32-bit libraries) <code>MQ_INSTALLATION_PATH/java/lib64</code> (64-bit libraries)
HP-UX Linux (POWER, x86-64 and zSeries s390x platforms) Solaris (x86-64 and SPARC platforms)	<code>MQ_INSTALLATION_PATH/java/lib</code> (32-bit libraries) <code>MQ_INSTALLATION_PATH/java/lib64</code> (64-bit libraries)
Linux (x86 platform)	<code>MQ_INSTALLATION_PATH/java/lib</code>
Windows	<code>MQ_INSTALLATION_PATH\Java\lib</code> (32-bit libraries) <code>MQ_INSTALLATION_PATH\Java\lib64</code> (64-bit libraries)
 z/OS	<code>MQ_INSTALLATION_PATH/mqm/V8R0M0/java/lib</code> (32-bit and 64-bit libraries)

**Note:**

-     On AIX, HP-UX, Linux (Power platform), or Solaris, use either the 32-bit libraries or the 64-bit libraries. Use the 64-bit libraries only if you are running your application in a 64-bit Java virtual machine (JVM) on a 64-bit platform. Otherwise, use the 32-bit libraries.
-  On Windows, you can use the PATH environment variable to specify the location of the IBM MQ classes for Java libraries instead of specifying their location on the **java** command.
-  To use IBM MQ classes for Java in bindings mode on IBM i, ensure that the library QMQMJAVA is in your library list.
-  On z/OS, you can use either a 32-bit or 64-bit Java virtual machine (JVM). You do not have to specify which libraries to use; IBM MQ classes for Java can determine for itself which JNI libraries to load.

**Related concepts:**

Using IBM MQ classes for Java

Use this information to learn how to configure your system to run the sample applications to verify your IBM MQ classes for Java installation. You can also find out how to modify the procedures to run your own applications.

*Support for OSGi with IBM MQ classes for Java:*

OSGi provides a framework that supports the deployment of applications as bundles. Three OSGi bundles are supplied as part of the IBM MQ classes for Java.

OSGi provides a general purpose, secure, and managed Java framework, which supports the deployment of applications that come in the form of bundles. OSGi-compliant devices can download and install bundles, and remove them when they are no longer required. The framework manages the installation and updating of bundles in a dynamic and scalable fashion.

The IBM MQ classes for Java include the following OSGi bundles.

**com.ibm.mq.osgi.java\_version\_number.jar**

The JAR files to allow applications to use the IBM MQ classes for Java.

**com.ibm.mq.osgi.allclient\_version\_number.jar**

This JAR file allows applications to use both the IBM MQ classes for JMS and the IBM MQ classes for Java, and also includes the code to handle PCF messages.

### **com.ibm.mq.osgi.allclientprereqs\_<version\_number>.jar**

This JAR file provides the prerequisites for `com.ibm.mq.osgi.allclient_<version_number>.jar`.

where *version\_number* is the version number of IBM MQ that has been installed.

The bundles are installed into the `java/lib/OSGi` subdirectory of your IBM MQ installation, or the `java\lib\OSGi` folder on Windows.

From IBM MQ Version 8.0, use the bundles `com.ibm.mq.osgi.allclient_8.0.0.0.jar`, and `com.ibm.mq.osgi.allclientprereqs_8.0.0.0.jar` for any new applications. Using these bundles removes the restriction of not being able to run both IBM MQ classes for JMS and the IBM MQ classes for Java within the same OSGi framework. All other restrictions still apply however. For versions of IBM MQ before Version 8.0, the restriction of using either IBM MQ classes for JMS or IBM MQ classes for Java applies.

Nine other bundles are also installed into the `java/lib/OSGi` sub directory of your IBM MQ installation, or the `java\lib\OSGi` folder on Windows. These bundles are part of the IBM MQ classes for JMS, and must not be loaded into an OSGi runtime environment that has the IBM MQ classes for Java bundle loaded. If the IBM MQ classes for Java OSGi bundle is loaded into an OSGi runtime environment that also has the IBM MQ classes for JMS bundles loaded, errors as shown in the following example occur when applications using either the IBM MQ classes for Java bundle or the IBM MQ classes for JMS bundles are run:


```
java.lang.ClassCastException: com.ibm.mq.MQException incompatible with com.ibm.mq.MQException
```

The OSGi bundle for the IBM MQ classes for Java has been written to the OSGi Release 4 specification; it does not work in an OSGi Release 3 environment.

You must set your system path or library path correctly so that the OSGi runtime environment can find any required DLL files or shared libraries.

If you use the OSGi bundle for the IBM MQ classes for Java, channel exit classes written in Java are not supported because of an inherent problem in loading classes in a multiple class loader environment such as OSGi. A user bundle can be aware of the IBM MQ classes for Java bundle, but the IBM MQ classes for Java bundle is not aware of any user bundle. As a result, the class loader used in an IBM MQ classes for Java bundle cannot load a channel exit class that is in a user bundle.

For more information about OSGi, see the OSGi alliance website.

*Installation of IBM MQ classes for Java on z/OS:* 

On z/OS, the STEPLIB used at runtime must contain the IBM MQ SCSQAUTH and SCSQANLE libraries.

From UNIX and Linux System Services, you can add these libraries by using a line in your `.profile` as shown in the following example, replacing `thlqual` with the high level data set qualifier that you chose when installing IBM MQ:

```
export STEPLIB=thlqual.SCSQAUTH:thlqual.SCSQANLE:$STEPLIB
```

In other environments, you typically need to edit the startup JCL to include SCSQAUTH on the STEPLIB concatenation:



```
STEPLIB DD DSN=thlqual.SCSQAUTH,DISP=SHR
 DD DSN=thlqual.SCSQANLE,DISP=SHR
```



The IBM MQ classes for Java configuration file:

An IBM MQ classes for Java configuration file specifies properties that are used to configure IBM MQ classes for Java.

The format of an IBM MQ classes for Java configuration file is that of a standard Java properties file.

  From IBM MQ Version 9.0.3 and IBM MQ Version 9.0.0, Fix Pack 2, a sample configuration file, `mqjava.config`, is supplied in the `bin` subdirectory of the IBM MQ classes for Java installation directory. This file documents all the supported properties and their default values.

**Note:** The sample configuration file is overwritten when the IBM MQ installation is upgraded to a future Fix Pack. Therefore, it is recommended that you make a copy of the sample configuration file for use with your applications.

You can choose the name and location of an IBM MQ classes for Java configuration file. When you start your application, use a **java** command with the following format:

```
java -Dcom.ibm.msg.client.config.location=config_file_url application_name
```

In the command, *config\_file\_url* is a uniform resource locator (URL) that specifies the name and location of the IBM MQ classes for Java configuration file. URLs of the following types are supported: `http`, `file`, `ftp`, and `jar`.

The following example shows a **java** command:

```
java -Dcom.ibm.msg.client.config.location=file:/D:/mydir/mqjava.config MyAppClass
```

This command identifies the IBM MQ classes for Java configuration file as the file `D:\mydir\mqjava.config` on the local Windows system.

When an application starts, IBM MQ classes for Java reads the contents of the configuration file and stores the specified properties in an internal property store. If the **java** command does not identify a configuration file, or if the configuration file cannot be found, IBM MQ classes for Java uses the default values for all the properties. If required, you can override any property in the configuration file by specifying it as a system property on the **java** command.

An IBM MQ classes for Java configuration file can be used with any of the supported transports between an application and a queue manager or broker.

### Overriding properties specified in an IBM MQ MQI client configuration file

An IBM MQ MQI client configuration file can also specify properties that are used to configure IBM MQ classes for Java. However, properties that are specified in an IBM MQ MQI client configuration file apply only when an application connects to a queue manager in client mode.

If required, you can override any attribute in an IBM MQ MQI client configuration file by specifying it as a property in an IBM MQ classes for Java configuration file. To override an attribute in an IBM MQ MQI client configuration file, use an entry with the following format in the IBM MQ classes for Java configuration file:

```
com.ibm.mq.cfg.stanza.propName=propValue
```

The variables in the entry have the following meanings:

**stanza** The name of the stanza in the IBM MQ MQI client configuration file that contains the attribute.

**propName**

The name of the attribute as specified in the IBM MQ MQI client configuration file.

### *propValue*

The value of the property that overrides the value of the attribute that is specified in the IBM MQ MQI client configuration file.

Alternatively, you can override an attribute in an IBM MQ MQI client configuration file by specifying the property as a system property on the **java** command. Use the preceding format to specify the property as a system property.

Only the following attributes in an IBM MQ MQI client configuration file are relevant to IBM MQ classes for Java. If you specify or override other attributes, it has no effect. Specifically, note that the `ChannelDefinitionFile` and `ChannelDefinitionDirectory` in the `CHANNELS` stanza of the client configuration file are not used. See "Using a client channel definition table with IBM MQ classes for Java" on page 1047 for details of how to use the CCDT with the IBM MQ classes for Java.

Table 152. Which stanza of the client configuration file contains which attribute

Stanza	Attribute
CHANNELS stanza of the client configuration file	Put1DefaultAlwaysSync
CHANNELS stanza of the client configuration file	PasswordProtection
ClientExitPath stanza of the client configuration file	ExitsDefaultPath
ClientExitPath stanza of the client configuration file	ExitsDefaultPath64
ClientExitPath stanza of the client configuration file	JavaExitsClasspath
JMQI stanza of the client configuration file	useMQCSPauthentication
MessageBuffer stanza of the client configuration file	MaximumSize
MessageBuffer stanza of the client configuration file	PurgeTime
MessageBuffer stanza of the client configuration file	UpdatePercentage
TCP stanza of the client configuration file	ClntRcvBuffSize
TCP stanza of the client configuration file	ClntSndBuffSize
TCP stanza of the client configuration file	Connect_Timeout
TCP stanza of the client configuration file	KeepAlive

For more information on the IBM MQ MQI client configuration, see *Configuring a client using a configuration file*.

### **Related information:**

Tracing IBM MQ classes for Java applications

*Java Standard Environment Trace stanza:*

Use the Java Standard Environment Trace Settings stanza to configure the IBM MQ classes for Java trace facility.

**com.ibm.msg.client.commonservices.trace.outputName = *traceOutputName***

*traceOutputName* is the directory and file name to which trace output is sent.

*traceOutputName* defaults to a file named `mqjms_PID.trc` in the current working directory where *PID* is the current process ID. If a process ID is unavailable, a random number is generated and prefixed with the letter `f`. To include the process ID in a file name you specify, use the string `%PID%`.

If you specify an alternative directory, it must exist, and you must have write permission for this directory. If you do not have write permission, the trace output is written to `System.err`.

**com.ibm.msg.client.commonservices.trace.include = *includeList***

*includeList* is a list of packages and classes that are traced, or the special values `ALL` or `NONE`.

Separate package or class names with a semicolon, ;. *includeList* defaults to ALL, and traces all packages and classes in IBM MQ classes for Java.

**Note:** You can include a package but then exclude subpackages of that package. For example, if you include package a.b and exclude package a.b.x, the trace includes everything in a.b.y and a.b.z, but not a.b.x or a.b.x.1.

**com.ibm.msg.client.commonservices.trace.exclude = *excludeList***

*excludeList* is a list of packages and classes that are not traced, or the special values ALL or NONE.

Separate package or class names with a semicolon, ;. *excludeList* defaults to NONE, and therefore excludes no packages and classes in IBM MQ classes for JMS from being traced.

**Note:** You can exclude a package but then include subpackages of that package. For example, if you exclude package a.b and include package a.b.x, the trace includes everything in a.b.x and a.b.x.1, but not a.b.y or a.b.z.

Any package or class that is specified, at the same level, as both included and excluded is included.

**com.ibm.msg.client.commonservices.trace.maxBytes = *maxArrayBytes***

*maxArrayBytes* is the maximum number of bytes that are traced from any byte arrays.

If *maxArrayBytes* is set to a positive integer, it limits the number of bytes in a byte-array that are written out to the trace file. It truncates the byte array after writing *maxArrayBytes* out. Setting *maxArrayBytes* reduces the size of the resulting trace file, and reduces the effect of tracing on the performance of the application.

A value of 0 for this property means that none of the contents of any byte arrays are sent to the trace file.

The default value is -1, which removes any limit on the number of bytes in a byte array that are sent to the trace file.

**com.ibm.msg.client.commonservices.trace.limit = *maxTraceBytes***

*maxTraceBytes* is the maximum number of bytes that are written to a trace output file.

*maxTraceBytes* works with *traceCycles*. If the number of bytes of trace written is near to the limit, the file is closed, and a new trace output file is started.

A value of 0 means that a trace output file has zero length. The default value is -1, which means that the amount of data to be written to a trace output file is unlimited.

**com.ibm.msg.client.commonservices.trace.count = *traceCycles***

*traceCycles* is the number of trace output files to cycle through.

If the current trace output file reaches the limit specified by *maxTraceBytes*, the file is closed. Further trace output is written to the next trace output file in sequence. Each trace output file is distinguished by a numeric suffix appended to the file name. The current or most recent trace output file is mqjms.trc.0, the next most recent trace output file is mqjms.trc.1. Older trace files follow the same numbering pattern up to the limit.

The default value of *traceCycles* is 1. If *traceCycles* is 1, when the current trace output file reaches its maximum size, the file is closed and deleted. A new trace output file with the same name is started. Therefore, only one trace output file exists at a time.

**com.ibm.msg.client.commonservices.trace.parameter = *traceParameters***

*traceParameters* controls whether method parameters and return values are included in the trace.

*traceParameters* defaults to TRUE. If *traceParameters* is set to FALSE, only method signatures are traced.

**com.ibm.msg.client.commonservices.trace.startup = *startup***

There is an initialization phase of IBM MQ classes for Java during which resources are allocated. The main trace facility is initialized during the resource allocation phase.

If *startup* is set to TRUE, startup trace is used. Trace information is produced immediately and includes the setup of all components, including the trace facility itself. Startup trace information can be used to diagnose configuration problems. Startup trace information is always written to System.err.

*startup* defaults to FALSE.

*startup* is checked before initialization is complete. For this reason, only specify the property on the command line as a Java system property. Do not specify it in the IBM MQ classes for Java configuration file.

**com.ibm.msg.client.commonservices.trace.compress = *compressedTrace***

Set *compressedTrace* to TRUE to compress trace output.

The default value of *compressedTrace* is FALSE.

If *compressedTrace* is set to TRUE, trace output is compressed. The default trace output file name has the extension .trz. If compression is set to FALSE, the default value, the file has the extension .trc to indicate it is uncompressed. However if the file name for the trace output has been specified in *traceOutputName* that name is used instead; no suffix is applied to the file.

Compressed trace output is smaller than uncompressed. Because there is less I/O, it can be written out faster than uncompressed trace. Compressed tracing has less effect on the performance of IBM MQ classes for Java than uncompressed tracing.

If *maxTraceBytes* and *traceCycles* are set, multiple compressed trace files are created in place of multiple flat files.

If IBM MQ classes for Java ends in an uncontrolled manner, a compressed trace file might not be valid. For this reason, trace compression must only be used when IBM MQ classes for Java closes down in a controlled manner. Only use trace compression if the problems being investigated do not cause the JVM itself to stop unexpectedly. Do not use trace compression when diagnosing problems that can result in System.Halt() shutdowns or abnormal, uncontrolled JVM terminations.

**com.ibm.msg.client.commonservices.trace.level = *traceLevel***

*traceLevel* specifies a filtering level for the trace. The defined trace levels are as follows:

- TRACE\_NONE: 0
- TRACE\_EXCEPTION: 1
- TRACE\_WARNING: 3
- TRACE\_INFO: 6
- TRACE\_ENTRYEXIT: 8
- TRACE\_DATA: 9
- TRACE\_ALL: Integer.MAX\_VALUE

Each trace level includes all lower levels. For example, if trace level is set at TRACE\_INFO, then any trace point with a defined level of TRACE\_EXCEPTION, TRACE\_WARNING, or TRACE\_INFO is written to the trace. All other trace points are excluded.

**com.ibm.msg.client.commonservices.trace.standalone = *standaloneTrace***

*standaloneTrace* controls whether the IBM MQ classes for Java client tracing service is used in a WebSphere Application Server environment.

If *standaloneTrace* is set to TRUE, the IBM MQ classes for Java client tracing properties are used to determine the trace configuration.

If *standaloneTrace* is set to FALSE, and the IBM MQ classes for Java client is running in an WebSphere Application Server container, the WebSphere Application Server trace service is used. The trace information that is generated depends upon the trace settings of the application server.

The default value of *standaloneTrace* is FALSE.

*IBM MQ classes for Java and software management tools:*

Software management tools such as Apache Maven can be used with the IBM MQ classes for Java.

Many large development organizations use these tools to centrally manage repositories of third-party libraries.

The IBM MQ classes for Java are composed of a number of JAR files. When you are developing Java language applications by using this API, an installation of either an IBM MQ Server, IBM MQ Client, or IBM MQ Client SupportPac is required on the machine where the application is being developed.

If you want to use a software management tool and add the JAR files that make up the IBM MQ classes for Java to a centrally managed repository, the following points must be observed:

- A repository or container must be made available only to developers within your organization. Any distribution outside of the organization is not permitted.
- The repository needs to contain a complete and consistent set of JAR files from a single IBM MQ release or Fix Pack.
- You are responsible for updating the repository with any maintenance provided by IBM Support.

From IBM MQ Version 8.0, the `com.ibm.mq.allclient.jar` JAR file needs to be installed into the repository.

From IBM MQ Version 9.0, the Bouncy Castle security provider and CMS support JAR files are required. For more information, see “What is installed for IBM MQ classes for Java” on page 1026 and Support for non-IBM JREs.

### **Running IBM MQ classes for Java applications under the Java Security Manager:**

IBM MQ classes for Java can run with the Java Security Manager enabled. To successfully run applications with the Java Security Manager enabled, you must configure your Java virtual machine (JVM) with a suitable policy definition file.


The simplest way to create a suitable policy definition file is to change the policy file supplied with the Java runtime environment (JRE). On most systems, this file is stored in the path `lib/security/java.policy`, relative to your JRE directory. You can edit policy files either by using your preferred editor or by using the `policytool` program supplied with your JRE.

You must give authority to the `com.ibm.mq.jmqi.jar` file so that it can:

- Create sockets (in client mode)
- Load the native library (in bindings mode)
- Read various properties from the environment

The system property `os.name` must be available to the IBM MQ classes for Java when running under the Java Security Manager.

An IBM MQ queue manager can send notifications to connected clients requesting a controlled closure of conversations (connection handles), for example when the queue manager is being quiesced. If a thread within a Java client receives one of these notifications at the same time as another thread within the client requests a new conversation, a deadlock can occur as both threads need access to the internal "connectionsLock" on the `RemoteConnectionSpecification` object.

 From IBM MQ Version 9.0.0, Fix Pack 3 and IBM MQ Version 9.0.5, the deadlock within the IBM MQ Java client is fixed. If your Java application uses the Java Security Manager, you must add the following permission to the `java.security.policy` file used by the application, otherwise, exceptions will be thrown to the application:

```
permission java.lang.RuntimePermission "modifyThread";
```

This RuntimePermission is required by the client as part of managing the assignment and closure of multiplexed conversations over TCP/IP connections to queue managers.

### Example policy file entry

Here is an example of a policy file entry that allows IBM MQ classes for Java to run successfully under the default security manager. Replace the string `MQ_INSTALLATION_PATH` in this example with the location where IBM MQ classes for Java are installed on your system.

```
grant codeBase "file: MQ_INSTALLATION_PATH/java/lib/*" {
//We need access to these properties, mainly for tracing
permission java.util.PropertyPermission "user.name","read";
permission java.util.PropertyPermission "os.name","read";
permission java.util.PropertyPermission "user.dir","read";
permission java.util.PropertyPermission "line.separator","read";
permission java.util.PropertyPermission "path.separator","read";
permission java.util.PropertyPermission "file.separator","read";
permission java.util.PropertyPermission "com.ibm.msg.client.commonservices.log.*","read";
permission java.util.PropertyPermission "com.ibm.msg.client.commonservices.trace.*","read";
permission java.util.PropertyPermission "Diagnostics.Java.Errors.Destination.FileName","read";
permission java.util.PropertyPermission "com.ibm.mq.commonservices","read";
permission java.util.PropertyPermission "com.ibm.mq.cfg.*","read";

//Tracing - we need the ability to control java.util.logging
permission java.util.logging.LoggingPermission "control";
// And access to create the trace file and read the log file - assumed to be in the current directory
permission java.io.FilePermission "*","read,write";

// We'd like to set up an mBean to control trace
permission javax.management.MBeanServerPermission "createMBeanServer";
permission javax.management.MBeanPermission "*","*";

// We need to be able to read manifests etc from the jar files in the installation directory
permission java.io.FilePermission "MQ_INSTALLATION_PATH/java/lib/-","read";

//Required if mqclient.ini/mqs.ini configuration files are used
permission java.io.FilePermission "MQ_DATA_DIRECTORY/mqclient.ini","read";
permission java.io.FilePermission "MQ_DATA_DIRECTORY/mqs.ini","read";

//For the client transport type.
permission java.net.SocketPermission "*","connect,resolve";

//For the bindings transport type.
permission java.lang.RuntimePermission "loadLibrary.*";

//For applications that use CCDT tables (access to the CCDT AMQCLCHL.TAB)
permission java.io.FilePermission "MQ_DATA_DIRECTORY/qmgrs/QM_NAME/@ipcc/AMQCLCHL.TAB","read";

//For applications that use User Exits
permission java.io.FilePermission "MQ_DATA_DIRECTORY/exits/*","read";
permission java.io.FilePermission "MQ_DATA_DIRECTORY/exits64/*","read";
permission java.lang.RuntimePermission "createClassLoader";

//Required for the z/OS platform
permission java.util.PropertyPermission "com.ibm.vm.bitmode","read";

// Used by the internal ConnectionFactory implementation
permission java.lang.reflect.ReflectPermission "suppressAccessChecks";

// Used for controlled class loading
permission java.lang.RuntimePermission "setContextClassLoader";

// Used to default the Application name in Client mode connections
permission java.util.PropertyPermission "sun.java.command","read";
```

```
// Used by the IBM JSSE classes
permission java.util.PropertyPermission "com.ibm.crypto.provider.AESNITrace","read";

//Required to determine if an IBM Java Runtime is running in FIPS mode,
//and to modify the property values status as required.
permission java.util.PropertyPermission "com.ibm.jsse2.usefipsprovider","read,write";
permission java.util.PropertyPermission "com.ibm.jsse2.JSSEFIPS","read,write";
//Required if an IBM FIPS provider is to be used for SSL communication.
permission java.security.SecurityPermission "insertProvider.IBMJCEFIPS";

// Required for non-IBM Java Runtimes that establish secure client
// transport mode connections using mutual TLS authentication
permission java.util.PropertyPermission "javax.net.ssl.keyStore","read";
permission java.util.PropertyPermission "javax.net.ssl.keyStorePassword","read";
};
```

```
V 9.0.0.3 // Required for Java applications that use the Java Security Manager permission
java.lang.RuntimePermission "modifyThread";
```


This example of a policy file enables the IBM MQ classes for Java to work correctly under the security manager, but you might still need to enable your own code to run correctly before your applications work.

The sample code shipped with IBM MQ classes for Java has not been specifically enabled for use with the security manager; however the IVT tests run with this policy file and the default security manager in place.

### Running IBM MQ classes for Java applications under CICS Transaction Server:

An IBM MQ classes for Java application can be run as a transaction under CICS Transaction Server.

To run an IBM MQ classes for Java application as a transaction under CICS Transaction Server for z/OS, perform the following steps:

1. Define the application and transaction to CICS by using the supplied CEDA transaction.
2. Ensure that the IBM MQ CICS adapter is installed in your CICS system.  (See Using IBM MQ with CICS for details.)
3. Ensure that the JVM environment specified in CICS includes the appropriate CLASSPATH and LIBPATH entries.
4. Initiate the transaction by using any of your normal processes.

For more information on running CICS Java transactions, refer to your CICS system documentation.

## Using IBM MQ classes for Java

Use this information to learn how to configure your system to run the sample applications to verify your IBM MQ classes for Java installation. You can also find out how to modify the procedures to run your own applications.

The procedures depend on the connection option that you want to use. Follow the instructions in the section that is appropriate for your requirements.

Remember to check the IBM MQ product readme file for the latest information, or for more specific information about your environment. The latest version of the product readme file is available on the IBM MQ, WebSphere MQ, and MQSeries product readmes web page.

Before attempting to run an IBM MQ classes for Java application in bindings mode, make sure that you have configured IBM MQ as described in Configuring.

### Configuring your queue manager to accept client connections from IBM MQ classes for Java:

To configure your queue manager to accept incoming connection requests from clients, define and permit use of a server connection channel and start a listener program.

See “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470 for details.

### Verifying your IBM MQ classes for Java installation with the sample application:

An installation verification program, MQIVP, is supplied with IBM MQ classes for Java. You can use this program to test all the connection modes of IBM MQ classes for Java.

The program prompts for a number of choices and other data to determine which connection mode you want to verify. Use the following procedure to verify your installation:

1. If you are going to run the program in client mode, configure your queue manager as described in “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470 noting that the queue to be used is SYSTEM.DEFAULT.LOCAL.QUEUE.
2. If you are going to run the program in client mode, see also “Using IBM MQ classes for Java” on page 1021.  
Perform the remaining steps of this procedure on the system on which you are going to run the program.
3. Make sure that you have updated your CLASSPATH environment variable according to the instructions in “Environment variables relevant to IBM MQ classes for Java” on page 1029.
4. Change Directory to `MQ_INSTALLATION_PATH/mqm/samp/wmqjava/samples`, where `MQ_INSTALLATION_PATH` is the path to your IBM MQ installation. Then at the command prompt, enter:

```
java -Djava.library.path= library_path MQIVP
```

where *library\_path* is the path to the IBM MQ classes for Java libraries (see “IBM MQ classes for Java libraries” on page 1030 ).

At the prompt marked (1):

- To use a TCP/IP connection, enter an IBM MQ server host name.
- To use native connection (bindings mode), leave the field blank (do not enter a name).

The program tries to:

- 1. Connect to the queue manager
- 2. Open the queue SYSTEM.DEFAULT.LOCAL.QUEUE, put a message on the queue, get a message from the queue, and then close the queue
- 3. Disconnect from the queue manager





- 4. Return a message if the operations are successful

Here is an example of the prompts and responses you might see. The actual prompts and your responses depend on your IBM MQ network.

```
Please enter the IP address of the MQ server : ipaddress(1)
Please enter the port to connect to : (1414)(2)
Please enter the server connection channel name : channelname(2)
Please enter the queue manager name : qmname
Success: Connected to queue manager.
Success: Opened SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Put a message to SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Got a message from SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Closed SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Disconnected from queue manager
```

```
Tests complete -
SUCCESS: This MQ Transport is functioning correctly.
Press Enter to continue ...
```

**Note:**

1.  On z/OS, leave the field blank at prompt marked <sup>(1)</sup>.
2. If you choose server connection, you do not see the prompts marked <sup>(2)</sup>.
3.  On IBM i, you can only issue the java MQIVP command from QShell. Alternatively, you can run the application by using the CL command RUNJAVA CLASS(MQIVP).

**Solving IBM MQ classes for Java problems:**

Initially, run the installation verification program. You might also have to use the trace facility.

If a program does not complete successfully, run the installation verification program, and follow the advice given in the diagnostic messages. This program is described in “Verifying your IBM MQ classes for Java installation with the sample application” on page 1040.

If the problems continue and you need to contact the IBM service team, you might be asked to turn on the trace facility. Do this as shown in the following example.

To trace the MQIVP program:

- Create a *com.ibm.mq.commonservices* properties file (see Using *com.ibm.mq.commonservices*).
- Enter the following command:

```
java -Dcom.ibm.mq.commonservices=commonservices_properties_file java
-Djava.library.path= library_path MQIVP -trace
```

where:

- *commonservices\_properties\_file* is the path (including the filename) to the *com.ibm.mq.commonservices* properties file.
- *library\_path* is the path to the IBM MQ classes for Java libraries (see “IBM MQ classes for Java libraries” on page 1030).

For more information about how to use trace, see Tracing IBM MQ classes for Java applications.

## Java and JMS client connectivity to batch applications running on z/OS

z/OS V 9.0.4 MQ, Adv, VUE

A JMS, or IBM MQ classes for Java, application on z/OS can connect to a queue manager on z/OS, that has the **ADVCAP(ENABLED)** attribute, by using a client connection.

A value of **ADVCAP(ENABLED)** applies only to a z/OS queue manager, licensed as IBM MQ Advanced for z/OS, Value Unit Edition (see IBM MQ for z/OS product identifiers), running with **QMGRPROD** set to **ADVANCEDVUE**.

See **DISPLAY QMGR** for more information on **ADVCAP** and **START QMGR** for more information on **QMGRPROD**.

Note that batch is the only environment supported; there is no support for JMS for CICS or JMS for IMS.

An IBM MQ classes for JMS, or IBM MQ classes for Java, application on z/OS cannot connect, using client mode connection, to a queue manager that is not running on z/OS, or to a queue manager that does not have the **ADVCAP(ENABLED)** option.

If a JMS application attempts to connect using client mode, and the application is not allowed to do so, exception message **JMSFMQ0005** is issued.

If an IBM MQ classes for Java application on z/OS attempts to connect using client mode, and it is not allowed to do so, **MQRC\_ENVIRONMENT\_ERROR** is returned.

V 9.0.5

### Advanced Message Security (AMS) support

From IBM MQ Version 9.0.5, IBM MQ classes for JMS or IBM MQ classes for Java client applications can use AMS when connecting to IBM MQ Advanced for z/OS, Value Unit Edition queue managers on remote z/OS systems.

A new key store type, `jceracfks`, is supported in `keystore.conf` on z/OS only, where:

- The property name prefix is `jceracfks` and this name prefix is case insensitive.
- The key store is a RACF keyring.
- Passwords are not required, and will be ignored. This is because RACF keyrings do not use passwords.
- If you specify the provider, the provider must be `IBMJCE`.

When you use `jceracfks` with AMS, the key store must be in the form: `safkeyring://user/keyring`, where:

- `safkeyring` is a literal and this name is case insensitive
- `user` is the RACF user id that owns the keyring
- `keyring` is the name of the RACF keyring and the name of the keyring is case sensitive

The following example uses the standard AMS keyring for user `JOHNDOE`:

```
jceracfks.keystore=safkeyring://JOHNDOE/drq.ams.keyring
```

## Writing IBM MQ classes for Java applications

This collection of topics provides information to assist with writing Java applications to interact with IBM MQ systems.

To use IBM MQ classes for Java to access IBM MQ queues, you write Java applications that contain calls that put messages onto, and get messages from, IBM MQ queues. For details of individual classes, see IBM MQ classes for Java .

### The IBM MQ classes for Java interface

The procedural IBM MQ application programming interface uses verbs, which act on objects. The Java programming interface uses objects, which you act upon by calling methods.

The procedural IBM MQ application programming interface is built around verbs such as these:

```
MQBACK, MQBEGIN, MQCLOSE, MQCONN, MQDISC,
MQGET, MQINQ, MQOPEN, MQPUT, MQSET, MQSUB
```

These verbs all take, as a parameter, a handle to the IBM MQ object on which they are to operate. Your program consists of a set of IBM MQ objects, which you act upon by calling methods on those objects.

When you use the procedural interface, you disconnect from a queue manager by using the call MQDISC(*Hconn*, *CompCode*, *Reason*), where *Hconn* is a handle to the queue manager.

In the Java interface, the queue manager is represented by an object of class MQQueueManager. You disconnect from the queue manager by calling the disconnect() method on that class.

```
// declare an object of type queue manager
MQQueueManager queueManager=new MQQueueManager();
...
// do something...
...
// disconnect from the queue manager
queueManager.disconnect();
```

### IBM MQ classes for Java connection modes:

The way you program for IBM MQ classes for Java has some dependencies on the connection modes you want to use.

If you use client connections, there are a number of differences from the IBM MQ MQI client but it is conceptually similar. If you use bindings mode you can use fastpath bindings and can issue the MQBEGIN command. You specify which mode to use by setting variables in the MQEnvironment class.

*IBM MQ classes for Java client connections:*

When IBM MQ classes for Java is used as a client, it is like the IBM MQ MQI client, but has a number of differences.

If you are programming for *IBM MQ classes for Java* for use as a client, be aware of the following differences:

- It supports only TCP/IP.
- It does not read any IBM MQ environment variables at startup.
- Information that would be stored in a channel definition and in environment variables can be stored in a class called Environment. Alternatively, this information can be passed as parameters when the connection is made.
- Error and exception conditions are written to a log specified in the MQException class. The default error destination is the Java console.
- Only the following attributes in an IBM MQ client configuration file are relevant to IBM MQ classes for Java. If you specify other attributes, they are ineffective.

Stanza	Attribute
ClientExitPath stanza of the client configuration file	ExitsDefaultPath
ClientExitPath stanza of the client configuration file	ExitsDefaultPath64
ClientExitPath stanza of the client configuration file	JavaExitsClasspath
MessageBuffer stanza of the client configuration file	MaximumSize
MessageBuffer stanza of the client configuration file	PurgeTime
MessageBuffer stanza of the client configuration file	UpdatePercentage
TCP stanza of the client configuration file	ClntRcvBuffSize
TCP stanza of the client configuration file	ClntSndBuffSize
TCP stanza of the client configuration file	Connect_Timeout
TCP stanza of the client configuration file	KeepAlive

- If connecting to a queue manager that requires character data to be converted, then the V7 Java client is now capable of doing the conversion if queue manager is unable to do so. The client JVM must support the conversion between the CCSID of the client and that of the queue manager.
- Automatic client reconnect is not supported by IBM MQ classes for Java.

When used in client mode, *IBM MQ classes for Java* does not support the MQBEGIN call.

*IBM MQ classes for Java bindings mode:*

The bindings mode of IBM MQ classes for Java differs from the client mode in three main ways.

When used in bindings mode, IBM MQ classes for Java uses the Java Native Interface (JNI) to call directly into the existing queue manager API, rather than communicating through a network.

By default, applications that use the IBM MQ classes for Java in bindings mode connect to a queue manager using the *ConnectOption*, MQCNO\_STANDARD\_BINDINGS.

The IBM MQ classes for Java support the following *ConnectOptions*:

- MQCNO\_FASTPATH\_BINDING
- MQCNO\_STANDARD\_BINDING
- MQCNO\_SHARED\_BINDING

- MQCNO\_ISOLATED\_BINDING

For further information on *ConnectOptions*, see “Connecting to a queue manager using the MQCONN call” on page 94.

Bindings mode supports the MQBEGIN call to initiate global units of work that are coordinated by the queue manager, on all platforms apart from IBM MQ for IBM i and IBM MQ for z/OS.

Most of the parameters provided by the MQEnvironment class are not relevant to bindings mode and are ignored.

*Defining which IBM MQ classes for Java connection to use:*

The type of connection to use is determined by the setting of variables in the MQEnvironment class.

Two variables are used:

#### **MQEnvironment.properties**

The connection type is determined by the value associated with the key name CMQC.TRANSPORT\_PROPERTY. Possible values are as follows:

##### **CMQC.TRANSPORT\_MQSERIES\_BINDINGS**

Connect in bindings mode

##### **CMQC.TRANSPORT\_MQSERIES\_CLIENT**

Connect in client mode

##### **CMQC.TRANSPORT\_MQSERIES**

Connection mode is determined by the value of the *hostname* property

#### **MQEnvironment.hostname**

Set the value of this variable as follows:

- For client connections, set the value of this variable to the host name of the IBM MQ server to which you want to connect
- For bindings mode, do not set this variable, or set it to null

#### **Operations on queue managers:**

This collection of topics describes how to connect to, and disconnect from, a queue manager using IBM MQ classes for Java.

*Setting up the IBM MQ environment for IBM MQ classes for Java:*

For an application to connect to a queue manager in client mode, the application must specify the channel name, host name, and port number.

**Note:** The information in this topic is relevant only if your application connects to a queue manager in client mode. It is not relevant if it connects in bindings mode. See: “Connection modes for IBM MQ classes for JMS” on page 851

You can specify the channel name, host name, and port number in one of two ways: either as fields in the MQEnvironment class or as properties of the MQQueueManager object.

If you set fields in the MQEnvironment class, they apply to your whole application, except where they are overridden by a properties hash table. To specify the channel name and host name in MQEnvironment, use the following code:

```
MQEnvironment.hostname = "host.domain.com";
MQEnvironment.channel = "java.client.channel";
```

This is equivalent to setting an **MQSERVER** environment variable:

```
"java.client.channel/TCP/host.domain.com".
```

By default, the Java clients attempt to connect to an IBM MQ listener at port 1414. To specify a different port, use the following code:

```
MQEnvironment.port = nnnn;
```

where *nnnn* is the required port number

If you pass properties to a queue manager object at its creation, they apply only to that queue manager. Create entries in a Hashtable object with keys of **hostname**, **channel**, and, optionally, **port**, and with appropriate values. To use the default port, 1414, you can omit the **port** entry. Create the MQQueueManager object by using a constructor that accepts the properties hash table.

### Identifying a connection to the queue manager by setting an application name

An application can set a name that identifies its connection to the queue manager. This application name is shown by the **DISPLAY CONN MQSC/PCF** command (where the field is called **APPLTAG**) or in the IBM MQ Explorer Application Connections display (where the field is called **App name**).

Application names are limited to 28 characters, so longer names are truncated. If an application name is not specified, a default is provided. The default name is based on the invoking (main) class, but if this information is not available, the text WebSphere MQ Client for Java is used.

If the name of the invoking class is used, it is adjusted to fit by removing leading package names, if necessary. For example, if the invoking class is `com.example.MainApp`, the full name is used, but if the invoking class is `com.example.dictionaryAndThesaurus.multilingual.mainApp`, the name `multilingual.mainApp` is used, because it is the longest combination of class name and rightmost package name that fits into the available length.

If the class name itself is more than 28 characters long, it is truncated to fit. For example, `com.example.mainApplicationForSecondTestCase` becomes `mainApplicationForSecondTest`.

To set an application name in the MQEnvironment class, add the name to the MQEnvironment.properties hash table, with a key of **MQConstants.APPNAME\_PROPERTY**, by using the following code:

```
MQEnvironment.properties.put(MQConstants.APPNAME_PROPERTY, "my_application_name");
```

To set an application name in the properties hash table that is passed to the MQQueueManager constructor, add the name to the properties hash table with a key of **MQConstants.APPNAME\_PROPERTY**.

### Overriding properties specified in an IBM MQ client configuration file

An IBM MQ client configuration file can also specify properties that are used to configure IBM MQ classes for Java. However, properties specified in an IBM MQ MQI client configuration file apply only when an application connects to a queue manager in client mode.

If required, you can override any attribute in an IBM MQ configuration file in any of the following ways. The options are shown in order of precedence.

- Set a Java system property for the configuration property.
- Set the property in the MQEnvironment.properties map.
- On Java5 and later releases, set a system environment variable.

Only the following attributes in an IBM MQ client configuration file are relevant to IBM MQ classes for Java. If you specify or override other attributes, it has no effect.

Stanza	Attribute
ClientExitPath stanza of the client configuration file	ExitsDefaultPath
ClientExitPath stanza of the client configuration file	ExitsDefaultPath64
ClientExitPath stanza of the client configuration file	JavaExitsClasspath
MessageBuffer stanza of the client configuration file	MaximumSize
MessageBuffer stanza of the client configuration file	PurgeTime
MessageBuffer stanza of the client configuration file	UpdatePercentage
TCP stanza of the client configuration file	ClntRcvBufSize
TCP stanza of the client configuration file	ClntSndBufSize
TCP stanza of the client configuration file	Connect_Timeout
TCP stanza of the client configuration file	KeepAlive

*Connecting to a queue manager in IBM MQ classes for Java:*

Connect to a queue manager by creating a new instance of the MQQueueManager class. Disconnect from a queue manager by calling the disconnect() method.

You are now ready to connect to a queue manager by creating a new instance of the MQQueueManager class:

```
MQQueueManager queueManager = new MQQueueManager("qMgrName");
```

To disconnect from a queue manager, call the disconnect() method on the queue manager:

```
queueManager.disconnect();
```

If you call the disconnect method, all open queues and processes that you have accessed through that queue manager are closed. However, it is good programming practice to close these resources explicitly when you finish using them. To do this, use the close() method on the relevant objects.

The commit() and backout() methods on a queue manager are equivalent to the MQCMIT and MQBACK calls that are used with the procedural interface.

*Using a client channel definition table with IBM MQ classes for Java:*

An IBM MQ classes for Java client application can use client connection channel definitions stored in a client channel definition table (CCDT).

As an alternative to creating a client connection channel definition by setting certain fields and environment properties in the MQEnvironment class or passing them to an MQQueueManager in a properties hash table, an IBM MQ classes for Java client application can use client connection channel definitions that are stored in a client channel definition table. These definitions are created by IBM MQ Script (MQSC) commands or IBM MQ Programmable Command Format (PCF) commands, or using the IBM MQ Explorer .

When the application creates an MQQueueManager object, the IBM MQ classes for Java client searches the client channel definition table for a suitable client connection channel definition, and uses the channel definition to start an MQI channel. For more information about client channel definition tables and how to construct one, see Client channel definition table.

To use a client channel definition table, an application must first create a URL object. The URL object encapsulates a uniform resource locator (URL) that identifies the name and location of the file containing the client channel definition table and specifies how the file can be accessed.

For example, if the file `ccdt1.tab` contains a client channel definition table and is stored on the same system on which the application is running, the application can create a URL object in the following way:

```
java.net.URL chanTab1 = new URL("file:///home/admdata/ccdt1.tab");
```

As another example, suppose the file `ccdt2.tab` contains a client channel definition table and is stored on a system that is different from the one on which the application is running. If the file can be accessed using the FTP protocol, the application can create a URL object in the following way:

```
java.net.URL chanTab2 = new URL("ftp://ftp.server/admdata/ccdt2.tab");
```

After the application has created a URL object, the application can create an `MQueueManager` object using one of the constructors that takes a URL object as a parameter. Here is an example:

```
MQueueManager mars = new MQueueManager("MARS", chanTab2);
```

This statement causes the IBM MQ classes for Java client to access the client channel definition table identified by the URL object `chanTab2`, search the table for a suitable client connection channel definition, and then use the channel definition to start an MQI channel to the queue manager called MARS.

Note the following points that apply if an application uses a client channel definition table:

- When the application creates an `MQueueManager` object using a constructor that takes a URL object as a parameter, no channel name must be set in the `MQueueEnvironment` class, either as a field or as an environment property. If a channel name is set, the IBM MQ classes for Java client throws an `MQueueException`. The field or environment property specifying the channel name is considered to be set if its value is anything other than null, an empty string, or a string containing all blank characters.
- The `queueManagerName` parameter on the `MQueueManager` constructor can have one of the following values:
  - The name of a queue manager
  - An asterisk (\*) followed by the name of a queue manager group
  - An asterisk (\*)
  - Null, an empty string, or a string containing all blank characters

These are the same values that can be used for the `QMGrName` parameter on an `MQueueConn` call issued by a client application that is using Message Queue Interface (MQI). For more information about the meaning of these values, see “The Message Queue Interface overview” on page 75.

If your application uses connection pooling, see “Controlling the default connection pool in IBM MQ classes for Java” on page 1069.

- When the IBM MQ classes for Java client finds a suitable client connection channel definition in the client channel definition table, it uses only the information extracted from this channel definition to start an MQI channel. Any channel related fields or environment properties that the application might have set in the `MQueueEnvironment` class are ignored.

In particular, note the following points if you are using Transport Layer Security (TLS):

- An MQI channel uses TLS only if the channel definition extracted from the client channel definition table specifies the name of a CipherSpec supported by the IBM MQ classes for Java client.
- A client channel definition table also contains information about the location of Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs). The IBM MQ classes for Java client uses only this information to access LDAP servers that hold CRLs.
- A client channel definition table can also contain the location of an OCSP responder. IBM MQ classes for Java cannot use the OCSP information in a client channel definition table file. However, you can configure OCSP as described in the section Using Online Certificate Protocol

For more information about using TLS with a client channel definition table, see Specifying that an MQI channel uses TLS.

Note also the following points if you are using channel exits:



- An MQI channel uses the channel exits and associated user data specified by the channel definition extracted from the client channel definition table in preference to channel exits and data specified using other methods.
- A channel definition extracted from a client channel definition table can specify channel exits that are written in Java, C, or C++. For more information about how to write a channel exit in Java, see "Creating a channel exit in IBM MQ classes for Java" on page 1063. For more information about how to write a channel exit in other languages, see "Using channel exits not written in Java with IBM MQ classes for Java" on page 1066.

*Specifying a range of ports for IBM MQ classes for Java client connections:*

You can specify a port, or a range of ports, that an application can bind to in either of two ways.

When an IBM MQ classes for Java application attempts to connect to an IBM MQ queue manager in client mode, a firewall might allow only those connections that originate from specified ports or range of ports. In this situation, you can specify a port, or a range of ports, that the application can bind to. You can specify the port(s) in the following ways:

- You can set the `localAddressSetting` field in the `MQEnvironment` class. Here is an example:  

```
MQEnvironment.localAddressSetting = "192.0.2.0(2000,3000)";
```
- You can set the environment property `CMQC.LOCAL_ADDRESS_PROPERTY`. Here is an example:  

```
(MQEnvironment.properties).put(CMQC.LOCAL_ADDRESS_PROPERTY,
 "192.0.2.0(2000,3000)");
```
- When you can construct the `MQQueueManager` object, you can pass a properties hashtable containing a `LOCAL_ADDRESS_PROPERTY` with the value "192.0.2.0(2000,3000)"

In each of these examples, when the application later connects to a queue manager, the application binds to a local IP address and port number in the range 192.0.2.0(2000) to 192.0.2.0(3000).

In a system with more than one network interface, you can also use the `localAddressSetting` field, or the environment property `CMQC.LOCAL_ADDRESS_PROPERTY`, to specify which network interface must be used for a connection.

Connection errors might occur if you restrict the range of ports. If an error occurs, an `MQException` is thrown containing the IBM MQ reason code `MQRC_Q_MGR_NOT_AVAILABLE` and the following message:

```
Socket connection attempt refused due to LOCAL_ADDRESS_PROPERTY restrictions
```

An error might occur if all the ports in the specified range are in use, or if the specified IP address, host name, or port number is not valid (a negative port number, for example).

## Accessing queues, topics, and processes in IBM MQ classes for Java:

To access queues, topics, and processes, use methods of the `MQQueueManager` class. The MQOD (object descriptor structure) is collapsed into the parameters of these methods.

### Queues

To open a queue you can use the `accessQueue` method of the `MQQueueManager` class. For example, on a queue manager called `queueManager`, use the following code:

```
MQQueue queue = queueManager.accessQueue("qName", CMQC.MQOO_OUTPUT);
```

The `accessQueue` method returns a new object of class `MQQueue`.

When you have finished using the queue, use the `close()` method to close it, as in the following example:  
`queue.close();`

You can also create a queue by using the `MQQueue` constructor. The parameters are exactly the same as for the `accessQueue` method, with the addition of a queue manager parameter. For example:

```
MQQueue queue = new MQQueue(queueManager,
 "qName",
 CMQC.MQOO_OUTPUT,
 "qMgrName",
 "dynamicQName",
 "altUserID");
```

You can specify a number of options when you create queues. For details of these, see `Class.com.ibm.mq.MQQueue`. Constructing a queue object in this way enables you to write your own subclasses of `MQQueue`.

### Topics

Similarly, you can open a topic using the `accessTopic` method of the `MQQueueManager` class. For example, on a queue manager called `queueManager`, use the following code to create a subscriber and publisher:

```
MQTopic subscriber =
 queueManager.accessTopic("TOPICSTRING", "TOPICNAME",
 CMQC.MQTOPIC_OPEN_AS_SUBSCRIPTION, CMQC.MQSO_CREATE);
MQTopic publisher =
 queueManager.accessTopic("TOPICSTRING", "TOPICNAME",
 CMQC.MQTOPIC_OPEN_AS_PUBLICATION, CMQC.MQOO_OUTPUT);
```

When you have finished using the topic, use the `close()` method to close it.

You can also create a topic by using the `MQTopic` constructor. The parameters are exactly the same as for the `accessTopic` method, with the addition of a queue manager parameter. For example:

```
MQTopic subscriber = new
 MQTopic(queueManager, "TOPICSTRING", "TOPICNAME",
 CMQC.MQTOPIC_OPEN_AS_SUBSCRIPTION, CMQC.MQSO_CREATE);
```

You can specify a number of options when you create topics. For details of these, see `Class com.ibm.mq.MQTopic`. Constructing a topic object in this way enables you to write your own subclasses of `MQTopic`.

A topic must be opened either for publication or for subscription. The `MQQueueManager` class has eight `accessTopic` methods and the `Topic` class has eight constructors. In each case, four of these have a **destination** parameter and four have a **subscriptionName** parameter (including two that have both).

These can only be used to open the topic for subscriptions. The two remaining methods have an **openAs** parameter, and the topic can be opened for either publication or subscription depending on the value of the **openAs** parameter.

To create a topic as a durable subscriber use either an `accessTopic` method of the `MQQueueManager` class or an `MQTopic` constructor that accepts a subscription name and, in either case, set the `CMQC.MQSO_DURABLE` option.

## Processes

To access a process, use the `accessProcess` method of the `MQQueueManager`. For example, on a queue manager called `queueManager`, use the following code to create an `MQProcess` object:

```
MQProcess process =
queueManager.accessProcess("PROCESSNAME",
CMQC.MQOO_FAIL_IF QUIESCING);
```

To access a process, use the `accessProcess` method of the `MQQueueManager`.

The `accessProcess` method returns a new object of class `MQProcess`.

When you have finished using the process object, use the `close()` method to close it, as in the following example:

```
process.close();
```

You can also create a process by using the `MQProcess` constructor. The parameters are exactly the same as for the `accessProcess` method, with the addition of a queue manager parameter. For example:

```
MQProcess process =
 new MQProcess(queueManager,"PROCESSNAME",
 CMQC.MQOO_FAIL_IF QUIESCING);
```

Constructing a process object in this way enables you to write your own subclasses of `MQProcess`.

## Handling messages in IBM MQ classes for Java:

Messages are represented by the `MQMessage` class. You put and get messages using methods of the `MQDestination` class, which has subclasses of `MQQueue` and `MQTopic`.

Put messages onto queues or topics using the `put()` method of the `MQDestination` class. You get messages from queues or topics using the `get()` method of the `MQDestination` class. Unlike the procedural interface, where `MQPUT` and `MQGET` put and get arrays of bytes, the Java programming language puts and gets instances of the `MQMessage` class. The `MQMessage` class encapsulates the data buffer that contains the actual message data, together with all the `MQMD` (message descriptor) parameters and message properties that describe that message.

To build a new message, create a new instance of the `MQMessage` class, and use the `writeXXX` methods to put data into the message buffer.

When the new message instance is created, all the `MQMD` parameters are automatically set to their default values, as defined in Initial values and language declarations for `MQMD`. The `put()` method of `MQDestination` also takes an instance of the `MQPutMessageOptions` class as a parameter. This class represents the `MQPMO` structure. The following example creates a message and puts it onto a queue:

```
// Build a new message containing my age followed by my name
MQMessage myMessage = new MQMessage();
myMessage.writeInt(25);
```

```
String name = "Charlie Jordan";
myMessage.writeInt(name.length());
```

```

myMessage.writeBytes(name);

// Use the default put message options...
MQPutMessageOptions pmo = new MQPutMessageOptions();

// put the message
!queue.put(myMessage,pmo);

```

The `get()` method of `MQDestination` returns a new instance of `MQMessage`, which represents the message just taken from the queue. It also takes an instance of the `MQGetMessageOptions` class as a parameter. This class represents the MQGMO structure.

You do not need to specify a maximum message size, because the `get()` method automatically adjusts the size of its internal buffer to fit the incoming message. Use the `readXXX` methods of the `MQMessage` class to access the data in the returned message.

The following example shows how to get a message from a queue:

```

// Get a message from the queue
MQMessage theMessage = new MQMessage();
MQGetMessageOptions gmo = new MQGetMessageOptions();
queue.get(theMessage,gmo); // has default values

// Extract the message data
int age = theMessage.readInt();
int strLen = theMessage.readInt();
byte[] strData = new byte[strLen];
theMessage.readFully(strData,0,strLen);
String name = new String(strData,0);

```

You can alter the number format that the read and write methods use by setting the *encoding* member variable.

You can alter the character set to use for reading and writing strings by setting the *characterSet* member variable.

See “MQMessage class” on page 728 for more information.

**Note:** The `writeUTF()` method of `MQMessage` automatically encodes the length of the string as well as the Unicode bytes it contains. When your message will be read by another Java program (using `readUTF()`), this is the simplest way to send string information.

*Improving the performance of nonpersistent messages in IBM MQ classes for Java:*

To improve performance when browsing messages or consuming nonpersistent messages from a client application, you can use *read ahead*. Client applications using `MQGET` or asynchronous consumption will benefit from the performance improvements when browsing messages or consuming nonpersistent messages.

For general information about the read ahead facility, see .

In IBM MQ classes for Java, you use the `CMQC.MQSO_READ_AHEAD` and `CMQC.MQSO_NO_READ_AHEAD` properties of an `MQQueue` or `MQTopic` object to determine whether message consumers and queue browsers are allowed to use read ahead on that object.

*Putting messages asynchronously using IBM MQ classes for Java:*

To put a message asynchronously, set `MQPMO_ASYNC_RESPONSE`.

You put messages onto queues or topics using the `put()` method of the `MQDestination` class. To put a message asynchronously, that is, allowing the operation to complete without waiting for a response from the queue manager, you can set `MQPMO_ASYNC_RESPONSE` in the options field of `MQPutMessageOptions`. To determine the success or failure of asynchronous puts, use the `MQQueueManager.getAsyncStatus` call.

### **Publish/subscribe in IBM MQ classes for Java:**

In IBM MQ classes for Java, the topic is represented by the `MQTopic` class, and you publish to it using the `MQTopic.put()` methods.

For general information about IBM MQ publish/subscribe, see [Publish/subscribe messaging](#).

### **Handling IBM MQ message headers with IBM MQ classes for Java:**

Java classes are provided representing different types of message header. Two helper classes are also provided.

#### **The MQHeader interface**

Header objects are described by the `MQHeader` interface, which provides general-purpose methods for accessing header fields and for reading and writing message content. Each header type has its own class that implements the `MQHeader` interface and adds getter and setter methods for individual fields. For example, the `MQRFH2` header type is represented by the `MQRFH2` class; the `MQDLH` header type by the `MQDLH` class, and so on. The header classes perform any necessary data conversion automatically, and can read or write data in any specified numeric encoding or character set (CCSID).

**Important:** The `MQRFH2` headers classes treat the message as a random access file, which means that the cursor must be positioned at the start of the message. Before using an internal message header class like `MQRFH`, `MQRFH2`, `MQCIH`, `MQDEAD`, `MQIIH` or `MQXMIT`, make sure that you update the message's cursor position to the correct location before passing the message to the class.

#### **Helper classes**

Two helper classes, `MQHeaderIterator` and `MQHeaderList`, assist with reading and decoding (parsing) the header content in messages:

- The `MQHeaderIterator` class works like a `java.util.Iterator`. For as long as there are more headers in the message, the `next()` method returns true, and the `nextHeader()` or `next()` method returns the next header object.
- The `MQHeaderList` works like a `java.util.List`. Like the `MQHeaderIterator`, it parses header content, but it also allows you to search for particular headers, add new headers, remove existing headers, update header fields and then write the header content back to a message. Alternatively, you can create an empty `MQHeaderList`, then populate it with header instances and write it to a message once or repeatedly.

The `MQHeaderIterator` and `MQHeaderList` classes use the information in the `MQHeaderRegistry` to know which IBM MQ header classes are associated with particular message types and formats. The `MQHeaderRegistry` is configured with knowledge of all current IBM MQ formats and header types and their implementation classes, and you can also register your own header types.

Support is provided for the following commonly used IBM MQ headers

- `MQRFH` - Rules and formatting header

- MQRFH2 - Like MQRFH, used to pass messages to and from a message broker belonging to IBM Integration Bus. Also used to contain message properties
- MQCIH - CICS Bridge
- MQDLH - Dead letter header
- MQIIH - IMS information header
- MQRMH - reference message header
- MQSAPH - SAP header
- MQWIH - Work information header
- MQXQH - Transmission Queue header
- MQDH - Distribution header
- MQEPH - Encapsulated PCF header

You can also define classes representing your own headers.

To use an MQHeaderIterator to get an RFH2 header, either set MQGMO\_PROPERTIES\_FORCE\_MQRFH2 in the GetMessageOptions, or set the queue property PROPCTL to FORCE.

*Printing all the headers in a message using IBM MQ classes for Java:*

In this example, an instance of MQHeaderIterator parses the headers in an MQMessage that has been received from a queue. The MQHeader objects returned from the nextHeader() method display their structure and contents when their toString method is invoked.

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQHeader;
import com.ibm.mq.headers.MQHeaderIterator;
...
MQMessage message = ... // Message received from a queue.
MQHeaderIterator it = new MQHeaderIterator (message);

while (it.hasNext ())
{
 MQHeader header = it.nextHeader ();

 System.out.println ("Header type " + header.type () + ": " + header);
}
```

*Skipping over the headers in a message using IBM MQ classes for Java:*

In this example, the skipHeaders() method of MQHeaderIterator positions the message read cursor immediately after the last header.

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQHeaderIterator;
...
MQMessage message = ... // Message received from a queue.
MQHeaderIterator it = new MQHeaderIterator (message);

it.skipHeaders ();
```

*Finding the reason code in a dead-letter message using IBM MQ classes for Java:*

In this example, the read method populates the MQDLH object by reading from the message. After the read operation, the message read cursor is positioned immediately after the MQDLH header content.

Messages on the queue manager's dead-letter queue are prefixed with a dead-letter header (MQDLH). To decide how to handle these messages - for example, to determine whether to retry or discard them - a dead-letter handling application must look at the reason code contained in the MQDLH.

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQDLH;
...
MQMessage message = ... // Message received from the dead-letter queue.
MQDLH dlh = new MQDLH ();

dlh.read (message);

System.out.println ("Reason: " + dlh.getReason ());
```

All header classes also provide a convenience constructor to initialize themselves directly from the message in a single step. So the code in this example could be simplified as follows:

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQDLH;
...
MQMessage message = ... // Message received from the dead-letter queue.
MQDLH dlh = new MQDLH (message);

System.out.println ("Reason: " + dlh.getReason ());
```

*Reading and removing the header from a dead-letter message using IBM MQ classes for Java:*

In this example, MQDLH is used to remove the header from a dead-letter message.

A dead-letter handling application will typically resubmit messages that have been rejected if their reason code indicates a transient error. Before resubmitting the message, it must remove the MQDLH header.

This example performs the following steps (see the comments in the example code):

1. The MQHeaderList reads the entire message, and each header encountered in the message becomes an item in the list.
2. Dead-letter messages contain an MQDLH as their first header, so this can be found in the first item of the header list. The MQDLH has already been populated from the message when the MQHeaderList is built, so there is no need to invoke its read method.
3. The reason code is extracted using the getReason() method provided by the MQDLH class.
4. The reason code has been inspected, and indicates that it is appropriate to resubmit the message. The MQDLH is removed using the MQHeaderList remove() method.
5. The MQHeaderList writes its remaining content to a new message object. The new message now contains everything in the original message except the MQDLH and can be written to a queue. The **true** argument to the constructor and to the write method indicates that the message body is to be held within the MQHeaderList, and written out again.
6. The format field in the message descriptor of the new message now contains the value that was previously in the MQDLH format field. The message data matches the numeric encoding and CCSID set in the message descriptor.

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQDLH;
import com.ibm.mq.headers.MQHeaderList;
...
MQMessage message = ... // Message received from the dead-letter queue.
MQHeaderList list = new MQHeaderList (message, true); // Step 1.
```

```

MQDLH dlh = (MQDLH) list.get (0); // Step 2.
int reason = dlh.getReason (); // Step 3.
...
list.remove (dlh); // Step 4.

MQMessage newMessage = new MQMessage ();

list.write (newMessage, true); // Step 5.
newMessage.format = list.getFormat (); // Step 6.

```

*Printing the content of a message using IBM MQ classes for Java:*

This example uses MQHeaderList to print out the content of a message, including its headers.

The output contains a view of all the header contents as well as the body of the message. The MQHeaderList class decodes all the headers in one go, whereas the MQHeaderIterator steps through them one at a time under application control. You might use this technique to provide a simple debugging tool when writing WebSphere MQ applications.

```

import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQHeaderList;
...
MQMessage message = ... // Message received from a queue.

System.out.println (new MQHeaderList (message, true));

```

This example also prints out the message descriptor fields, using the MQMD class. The copyFrom() method of the com.ibm.mq.headers.MQMD class populates the header object from the message descriptor fields of the MQMessage rather than by reading the message body.

```

import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQMD;
import com.ibm.mq.headers.MQHeaderList;
...
MQMessage message = ...
MQMD md = new MQMD ();
...
md.copyFrom (message);
System.out.println (md + "\n" + new MQHeaderList (message, true));

```

*Finding a specific type of header in a message using IBM MQ classes for Java:*

This example uses the indexOf(String) method of MQHeaderList to find an MQRFH2 header in a message, if one is present.

```

import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQHeaderList;
import com.ibm.mq.headers.MQRFH2;
...
MQMessage message = ...
MQHeaderList list = new MQHeaderList (message);
int index = list.indexOf ("MQRFH2");

if (index >= 0)
{
 MQRFH2 rfh = (MQRFH2) list.get (index);
 ...
}

```



*Analyzing an MQRFH2 header using IBM MQ classes for Java:*

This example shows how to access a known field value in a named folder, using the MQRFH2 class.

The MQRFH2 class provides a number of ways to access not only the fields in the fixed part of the structure, but also the XML-encoded folder contents that are carried within the NameValueData field. This example shows how to access a known field value in a named folder - in this instance, the Rto field in the jms folder, which represents the reply queue name in an MQ JMS message.

```
MQRFH2 rfh = ...
String value = rfh.getStringFieldValue ("jms", "Rto");
```

To discover the contents of an MQRFH2 (as opposed to requesting specific fields directly), you can use the getFolders method to return a list of MQRFH2.Element, which represents the structure of a folder that could contain fields and other folders. Setting a field or folder to null removes it from the MQRFH2. When you manipulate the NameValueData folder contents in this way, the StrucLength field is automatically updated accordingly.

*Reading and writing byte streams other than MQMessage objects using IBM MQ classes for Java:*

These examples use the header classes to parse and manipulate IBM MQ header content when the data source is not an MQMessage object.

You can use the header classes to parse and manipulate IBM MQ header content even when the data source is something other than an MQMessage object. The MQHeader interface implemented by every header class provides the methods int read (java.io.DataInput message, int encoding, int characterSet) and int write (java.io.DataOutput message, int encoding, int characterSet). The com.ibm.mq.MQMessage class implements the java.io.DataInput and java.io.DataOutput interfaces. This means that you can use the two MQHeader methods to read and write MQMessage content, overriding the encoding and CCSID specified in the message descriptor. This is useful for messages that contain a chain of headers in different encodings.

You can also obtain DataInput and DataOutput objects from other data streams, for example file or socket streams, or byte arrays carried in JMS messages. The java.io.DataInputStream classes implement DataInput and the java.io.DataOutputStream classes implement DataOutput. This example reads IBM MQ header content from a byte array:

```
import java.io.*;
import com.ibm.mq.headers.*;
...
byte [] bytes = ...
DataInput in = new DataInputStream (new ByteArrayInputStream (bytes));
MQHeaderIterator it = new MQHeaderIterator (in, CMQC.MQENC_NATIVE,
 CMQC.MQCCSI_DEFAULT);
```

The line starting MQHeaderIterator could be replaced with

```
MQDLH dlh = new MQDLH (in, CMQC.MQENC_NATIVE, CMQC.MQCCSI_DEFAULT);
// or any other header type
```

This example writes to a byte array using a DataOutputStream:

```
MQHeader header = ... // Could be any header type
ByteArrayOutputStream out = new ByteArrayOutputStream ();

header.write (new DataOutputStream (out), CMQC.MQENC_NATIVE, CMQC.MQCCSI_DEFAULT);
byte [] bytes = out.toByteArray ();
```

When you work with streams in this way, be careful to use the correct values for the encoding and characterSet arguments. When reading headers, specify the encoding and CCSID with which the byte content was originally written. When writing headers, specify the encoding and CCSID that you want to

produce. The data conversion is performed automatically by the header classes.

*Creating classes for new header types using IBM MQ classes for Java:*

You can create Java classes for header types not supplied with IBM MQ classes for Java.

To add a Java class representing a new header type that you can use in the same way as any header class supplied with IBM MQ classes for Java, you create a class that implements the MQHeader interface. The simplest approach is to extend the `com.ibm.mq.headers.impl.Header` class. This example produces a fully-functional class representing the MQTM header structure. You do not have to add individual getter and setter methods for each field, but it is a useful convenience for users of the header class. The generic `getValue` and `setValue` methods that take a string for the field name will work for all fields defined in the header type. The inherited `read`, `write` and `size` methods will enable instances of the new header type to be read and written and will calculate the header size correctly based upon its field definition. The type definition is created just once, however many instances of this header class are created. To make the new header definition available for decoding using the `MQHeaderIterator` or `MQHeaderList` classes, you would register it using the `MQHeaderRegistry`. Note however that the MQTM header class is in fact already provided in this package and registered in the default registry.

```
import com.ibm.mq.headers.impl.Header;
import com.ibm.mq.headers.impl.HeaderField;
import com.ibm.mq.headers.CMQC;

public class MQTM extends Header {
 final static HeaderType TYPE = new HeaderType ("MQTM");
 final static HeaderField StrucId = TYPE.addMQChar ("StrucId", CMQC.MQTM_STRUC_ID);
 final static HeaderField Version = TYPE.addMQLong ("Version", CMQC.MQTM_VERSION_1);
 final static HeaderField QName = TYPE.addMQChar ("QName", CMQC.MQ_Q_NAME_LENGTH);
 final static HeaderField ProcessName = TYPE.addMQChar ("ProcessName",
 CMQC.MQ_PROCESS_NAME_LENGTH);
 final static HeaderField TriggerData = TYPE.addMQChar ("TriggerData",
 CMQC.MQ_TRIGGER_DATA_LENGTH);
 final static HeaderField ApplType = TYPE.addMQLong ("ApplType");
 final static HeaderField ApplId = TYPE.addMQChar ("ApplId", 256);
 final static HeaderField EnvData = TYPE.addMQChar ("EnvData", 128);
 final static HeaderField UserData = TYPE.addMQChar ("UserData", 128);

 protected MQTM (HeaderType type){
 super (type);
 }
 public String getStrucId () {
 return getStringValue (StrucId);
 }
 public int getVersion () {
 return getIntValue (Version);
 }
 public String getQName () {
 return getStringValue (QName);
 }
 public void setQName (String value) {
 setStringValue (QName, value);
 }
 // ...Add convenience getters and setters for remaining fields in the same way.
}
```

## Handling PCF messages with IBM MQ classes for Java:

Java classes are provided to create and parse PCF-structured messages, and to facilitate sending PCF requests and collecting PCF responses.

Classes PCFMessage & MQCFGR represent arrays of PCF parameter structures. They provide convenience methods for adding and retrieving PCF parameters.

PCF parameter structures are represented by the classes MQCFH, MQCFIN, MQCFIN64, MQCFST, MQCFBS, MQCFIL, MQCFIL64, MQCFSL, and MQCFGR. These share basic operational interfaces:

- Methods to read and write message content: read (), write (), and size ()
- Methods to manipulate parameters: getValue (), setValue (), getParameter () and others
- The enumerator method .nextParameter (), which parses PCF content in an MQMessage

The PCF filter parameter is used in inquire commands to provide a filter function. It is encapsulated in the following classes:

- MQCFIF - integer filter
- MQCFSF - string filter
- MQCFBF - byte filter

Two agent classes, PCFAgent and PCFMessageAgent are provided to manage the connection to a Queue Manager, the command server queue, and an associated response queue. PCFMessageAgent extends PCFAgent and should normally be used in preference to it. The PCFMessageAgent class converts the received MQMessages and passes them back to the caller as a PCFMessage array. PCFAgent returns an array of MQMessages, which you have to parse before use.

## Handling message properties in IBM MQ classes for Java:

Function calls to process message handles have no equivalent in IBM MQ classes for Java. To set, return, or delete message handle properties, use methods of the MQMessage class.

For general information about message properties, see “Property names” on page 19.

In IBM MQ classes for Java access to messages is through the MQMessage class. Message handles are therefore not provided in the Java environment and there is no equivalent to the IBM MQ function calls MQCRTMH, MQDLTMH, MQMHBUF, and MQBUFMH

To set message handle properties in the procedural interface, you use the call MQSETMP. In IBM MQ classes for Java, use the appropriate method of the MQMessage class:

- setBooleanProperty
- setByteProperty
- setBytesProperty
- setShortProperty
- setIntProperty
- setInt2Property
- setInt4Property
- setInt8Property
- setLongProperty
- setFloatProperty
- setDoubleProperty
- setStringProperty

- setObjectProperty

These are sometimes referred to collectively as the *set\*property* methods.

To return the value of message handle properties in the procedural interface, you use the call MQINQMP. In IBM MQ classes for Java, use the appropriate method of the MQMessage class:

- getBooleanProperty
- getByteProperty
- getBytesProperty
- getShortProperty
- getIntProperty
- getInt2Property
- getInt4Property
- getInt8Property
- getLongProperty
- getFloatProperty
- getDoubleProperty
- getStringProperty
- getObjectProperty

These are sometimes referred to collectively as the *get\*property* methods.

To delete the value of message handle properties in the procedural interface, you use the call MQDLTMP. In IBM MQ classes for Java, use the deleteProperty method of the MQMessage class.

### Handling errors in IBM MQ classes for Java:

Handle errors arising from IBM MQ classes for Java using Java try and catch blocks.

Methods in the Java interface do not return a completion code and reason code. Instead, they throw an exception whenever the completion code and reason code resulting from an IBM MQ call are not both zero. This simplifies the program logic so that you do not have to check the return codes after each call to IBM MQ. You can decide at which points in your program you want to deal with the possibility of failure. At these points, you can surround your code with try and catch blocks, as in the following example:

```
try {
 myQueue.put(messageA,putMessageOptionsA);
 myQueue.put(messageB,putMessageOptionsB);
}
catch (MQException ex) {
 // This block of code is only executed if one of
 // the two put methods gave rise to a non-zero
 // completion code or reason code.
 System.out.println("An error occurred during the put operation:" +
 "CC = " + ex.completionCode +
 "RC = " + ex.reasonCode);
 System.out.println("Cause exception:" + ex.getCause());
}
```

The IBM MQ call reason codes reported back in Java exceptions for z/OS are documented in Reason codes for z/OS and Reason codes for all other platforms.

Exceptions that are thrown while an IBM MQ classes for Java application is running are also written to the log. However, an application can call the MQException.logExclude() method to prevent exceptions associated with a specific reason code from being logged. You might want to do this in situations where you expect many exceptions associated with a specific reason code to be thrown, and you do not want

the log to be filled with these exceptions. For example, if your application attempts to get a message from a queue each time it iterates around a loop and, for most of these attempts, you expect no suitable message to be on the queue, you might want to prevent exceptions associated with the reason code MQRC\_NO\_MSG\_AVAILABLE from being logged. If an application has previously prevented exceptions associated with a specific reason code from being logged, it can allow these exceptions to be logged again by calling the method MQException.logInclude().

Sometimes the reason code does not convey all details associated with the error. For each exception that is thrown, an application should check the linked exception. The linked exception itself might have another linked exception, and so the linked exceptions form a chain leading back to the original underlying problem. A linked exception is implemented by using the chained exception mechanism of the java.lang.Throwable class, and an application obtains a linked exception by calling the Throwable.getCause() method. From an exception that is an instance of MQException, MQException.getCause() retrieves the underlying instance of com.ibm.mq.jmqi.JmqiException, and getCause from this exception retrieves the underlying java.lang.Exception that caused the error.

### Getting and setting attribute values in IBM MQ classes for Java:

getXXX() and setXXX() methods are provided for many common attributes. Others can be accessed using the generic inquire() and set() methods.

For many of the common attributes, the classes MQManagedObject, MQDestination, MQQueue, MQTopic, MQProcess, and MQQueueManager contain getXXX() and setXXX() methods. These methods allow you to get and set their attribute values. Note that for MQDestination, MQQueue, and MQTopic, the methods work only if you specify the appropriate inquire and set flags when you open the object.

For less common attributes, the MQQueueManager, MQDestination, MQQueue, MQTopic, and MQProcess classes all inherit from a class called MQManagedObject. This class defines the inquire() and set() interfaces.

When you create a new queue manager object by using the *new* operator, it is automatically opened for inquire. When you use the accessProcess() method to access a process object, that object is automatically opened for inquire. When you use the accessQueue() method to access a queue object, that object is not automatically opened for either inquire or set operations. This is because adding these options automatically can cause problems with some types of remote queues. To use the inquire, set, getXXX, and setXXX methods on a queue, you must specify the appropriate inquire and set flags in the openOptions parameter of the accessQueue() method. The same is true for destination and topic objects.

The inquire and set methods take three parameters:

- selectors array
- intAttrs array
- charAttrs array

You do not need the SelectorCount, IntAttrCount, and CharAttrLength parameters that are found in MQINQ, because the length of an array in Java is always known. The following example shows how to make an inquiry on a queue:

```
// inquire on a queue
final static int MQIA_DEF_PRIORITY = 6;
final static int MQCA_Q_DESC = 2013;
final static int MQ_Q_DESC_LENGTH = 64;

int[] selectors = new int[2];
int[] intAttrs = new int[1];
byte[] charAttrs = new byte[MQ_Q_DESC_LENGTH]

selectors[0] = MQIA_DEF_PRIORITY;
selectors[1] = MQCA_Q_DESC;
```

```
queue.inquire(selectors,intAttrs,charAttrs);

System.out.println("Default Priority = " + intAttrs[0]);
System.out.println("Description : " + new String(charAttrs,0));
```

### **Multithreaded programs in Java:**

The Java runtime environment is inherently multithreaded. IBM MQ classes for Java allows a queue manager object to be shared by multiple threads but ensures that all access to the target queue manager is synchronized.

Multithreaded programs are hard to avoid in Java. Consider a simple program that connects to a queue manager and opens a queue at startup. The program displays a single button on the screen. When a user clicks that button, the program fetches a message from the queue.

The Java runtime environment is inherently multithreaded. Therefore, your application initialization occurs in one thread, and the code that executes in response to the button press executes in a separate thread (the user interface thread).

With the C based IBM MQ MQI client, this would cause a problem, because there are limitations to the sharing of handles by multiple threads. IBM MQ classes for Java relaxes this constraint, allowing a queue manager object (and its associated queue, topic and process objects) to be shared by multiple threads.

The implementation of IBM MQ classes for Java ensures that, for a particular connection (MQQueueManager object instance), all access to the target IBM MQ queue manager is synchronized. A thread that wants to issue a call to a queue manager is blocked until all other calls in progress for that connection are complete. If you require simultaneous access to the same queue manager from multiple threads within your program, create a new MQQueueManager object for each thread that requires concurrent access. (This is equivalent to issuing a separate MQCONN call for each thread.)

**Note:** Instances of the class `com.ibm.mq.MQGetMessageOptions` must not be shared between threads which are requesting messages concurrently. Instances of this class are updated with data during the corresponding MQGET request, which can result in unexpected consequences when multiple threads are operating concurrently on the same instance of the object.

### **Using channel exits in IBM MQ classes for Java:**

An overview of how to use channel exits in an application using the IBM MQ classes for Java.

The following topics describe how to write a channel exit in Java, how to assign it, and how to pass data to it. They then describe how to use channel exits written in C and how to use a sequence of channel exits.

Your application must have the correct security permission to load the channel exit class.

*Creating a channel exit in IBM MQ classes for Java:*

You can provide your own channel exits by defining a Java class that implements an appropriate interface.

To implement an exit, you define a new Java class that implements the appropriate interface. Three exit interfaces are defined in the `com.ibm.mq.exits` package:

- `WMQSendExit`
- `WMQReceiveExit`
- `WMQSecurityExit`

**Note:** Channel exits are supported for client connections only; they are not supported for bindings connections. You cannot use a Java channel exit outside IBM MQ classes for Java, for example if you are using a client application written in C.

Any TLS encryption defined for a connection is performed *after* send and security exits have been invoked. Similarly, decryption is performed *before* receive and security exits are invoked.

The following sample defines a class that implements all three interfaces:

```
public class MyMQExits implements
WMQSendExit, WMQReceiveExit, WMQSecurityExit {
 // Default constructor
 public MyMQExits(){
 }
 // This method comes from the send exit interface
 public ByteBuffer channelSendExit(
MQCXP channelExitParms,
 MQCD channelDefinition,
 ByteBuffer agentBuffer)
 {
 // Fill in the body of the send exit here
 }
 // This method comes from the receive exit interface
 public ByteBuffer channelReceiveExit(
MQCXP channelExitParms,
 MQCD channelDefinition,
 ByteBuffer agentBuffer)
 {
 // Fill in the body of the receive exit here
 }
 // This method comes from the security exit interface
 public ByteBuffer channelSecurityExit(
MQCXP channelExitParms,
 MQCD channelDefinition,
 ByteBuffer agentBuffer)
 {
 // Fill in the body of the security exit here
 }
}
```

Each exit is passed an MQCXP object and an MQCD object. These objects represent the MQCXP and MQCD structures defined in the procedural interface.

Any exit class you write must have a constructor. This can be either the default constructor or one that takes a string argument. If it takes a string then the user data will be passed into the exit class when it is created. If the exit class contains both a default constructor and a single argument constructor, the single argument constructor has priority.

For the send and security exits, your exit code must return the data that you want to send to the server. For a receive exit, your exit code must return the modified data that you want IBM MQ to interpret.

The simplest possible exit body is:

```
{ return agentBuffer; }
```

Do not close the queue manager from within a channel exit.

### Using existing channel exit classes

In versions of IBM MQ earlier than 7.0, you would implement these exits using the interfaces MQSendExit, MQReceiveExit, and MQSecurityExit, as in the following example. This method remains valid, but the new method is preferred for improved functionality and performance.

```
public class MyMQExits implements MQSendExit, MQReceiveExit, MQSecurityExit {
 // Default constructor
 public MyMQExits(){
 }
 // This method comes from the send exit
 public byte[] sendExit(MQChannelExit channelExitParms,
 MQChannelDefinition channelDefParms,
 byte agentBuffer[])
 {
 // Fill in the body of the send exit here
 }
 // This method comes from the receive exit
 public byte[] receiveExit(MQChannelExit channelExitParms,
 MQChannelDefinition channelDefParms,
 byte agentBuffer[])
 {
 // Fill in the body of the receive exit here
 }
 // This method comes from the security exit
 public byte[] securityExit(MQChannelExit channelExitParms,
 MQChannelDefinition channelDefParms,
 byte agentBuffer[])
 {
 // Fill in the body of the security exit here
 }
}
```

*Assigning a channel exit in IBM MQ classes for Java:*

You can assign a channel exit using IBM MQ classes for Java.

There is no direct equivalent to the IBM MQ channel in IBM MQ classes for Java. Channel exits are assigned to an MQQueueManager. For example, having defined a class that implements the WMQSecurityExit interface, an application can use the security exit in one of four ways:

- By assigning an instance of the class to the MQEnvironment.channelSecurityExit field before creating an MQQueueManager object
- By setting the MQEnvironment.channelSecurityExit field to a string representing the security exit class before creating an MQQueueManager object
- By creating a key/value pair in the properties hashtable passed to MQQueueManager with a key of CMQC.SECURITY\_EXIT\_PROPERTY
- Using a client channel definition table (CCDT)

Any exit assigned by setting the MQEnvironment.channelSecurityExit field to a string, creating a key/value pair in the properties hashtable, or using a CCDT, must be written with a default constructor. An exit assigned as an instance of a class does not need a default constructor, depending on the application.



An application can use a send or a receive exit in a similar way. For example, the following code fragment shows you how to use the security, send, and receive exits that are implemented in the class `MyMQExits`, which was defined previously, using `MQEnvironment`:

```
MyMQExits myexits = new MyMQExits();
MQEnvironment.channelSecurityExit = myexits;
MQEnvironment.channelSendExit = myexits;
MQEnvironment.channelReceiveExit = myexits;
:
MQQueueManager jupiter = new MQQueueManager("JUPITER");
```

If more than one method is used to assign a channel exit, the order of precedence is as follows:

1. If the URL of a CCDT is passed to the `MQQueueManager`, the contents of the CCDT determine the channel exits to be used and any exit definitions in `MQEnvironment` or the properties hashtable are ignored.
2. If no CCDT URL is passed, exit definitions from `MQEnvironment` and the hashtable are merged
  - If the same exit type is defined in both `MQEnvironment` and the hashtable, the definition in the hashtable is used.
  - If equivalent old and new types of exit are specified (for example the `sendExit` field, which can only be used for the type of exit used in versions of IBM MQ earlier than Version 7.0, and the `channelSendExit` field, which can be used for any send exit), the new exit (`channelSendExit`) is used rather than the old exit.

If you have declared a channel exit as a string, you must enable IBM MQ to locate the channel exit program. You can do so in various ways, depending on the environment in which the application is running and on how the channel exit programs are packaged.

- For an application that is running in an application server, you must store the files in the directory shown in Table 153 or packaged in JAR files referenced by **`exitClasspath`**.
- For an application that is not running in an application server, the following rules apply:
  - If your channel exit classes are packaged in separate JAR files, these JAR files must be included in the **`exitClasspath`**.
  - If your channel exit classes are not packaged in JAR files, the class files can be stored in the directory shown in Table 153 or in any directory in the JVM system class path or **`exitClasspath`**.

The **`exitClasspath`** property can be specified in four ways. In order of priority, these ways are as follows:

1. The system property `com.ibm.mq.exitClasspath` (defined on the command line using the `-D` option)
2. The `exitPath` stanza of the `mqclient.ini` file
3. A hashtable entry with the key `CMQC.EXIT_CLASSPATH_PROPERTY`
4. The `MQEnvironment` variable **`exitClasspath`**

Separate multiple paths using the `java.io.File.pathSeparator` character.

Table 153. The directory for channel exit programs

Platform	Directory
AIX , HP-UX, Linux , and Solaris	<code>/var/mqm/exits</code> (32-bit channel exit programs) <code>/var/mqm/exits64</code> (64-bit channel exit programs)
Windows	<code>install_data_dir\exits</code>

**Note:** `install_data_dir` is the directory that you chose for the IBM MQ data files during installation. The default directory is `C:\ProgramData\IBM\MQ`.

*Passing data to channel exits in IBM MQ classes for Java:*

You can pass data to channel exits and return data from channel exits to your application.

### **The agentBuffer parameter**

For a send exit, the *agentBuffer* parameter contains the data that is about to be sent. For a receive exit or a security exit, the *agentBuffer* parameter contains the data that has just been received. You do not need a length parameter, because the expression `agentBuffer.limit()` indicates the length of the array.

For the send and security exits, your exit code must return the data that you want to send to the server. For a receive exit, your exit code must return the modified data that you want IBM MQ to interpret.

The simplest possible exit body is:

```
{ return agentBuffer; }
```

Channel exits are called with a buffer that has a backing array. For best performance, the exit should return a buffer with a backing array.

### **User data**

If an application connects to a queue manager by setting `channelSecurityExit`, `channelSendExit`, or `channelReceiveExit`, 32 bytes of user data can be passed to the appropriate channel exit class when it is called, using the `channelSecurityExitUserData`, `channelSendExitUserData`, or `channelReceiveExitUserData` fields. This user data is available to the channel exit class but is refreshed each time the exit is called. Any changes made to the user data in the channel exit will therefore be lost. If you want to make persistent changes to data in a channel exit, use the MQCXP `exitUserArea`. Data in this field is maintained between invocations of the exit.

If the application sets `securityExit`, `sendExit`, or `receiveExit`, no user data can be passed to these channel exit classes.

If an application uses a client channel definition table (CCDT) to connect to a queue manager, any user data specified in a client connection channel definition is passed to channel exit classes when they are called. For more information about using a client channel definition table, see "Using a client channel definition table with IBM MQ classes for Java" on page 1047.

*Using channel exits not written in Java with IBM MQ classes for Java:*

How to use channel exit programs written in C from a Java application.

In IBM WebSphere MQ Version 7.0, you can specify the name of a channel exit program written in C as a String passed to the `channelSecurityExit`, `channelSendExit`, or `channelReceiveExit` fields in the `MQEnvironment` object or properties Hashtable. However, you cannot use a channel exit written in Java in an application written in another language.

Specify the exit program name in the format `library(function)` and ensure that the location of the exit program is specified as described in Path to exits.

### **Using external exit classes**

In versions of IBM MQ earlier than Version 7.0, three classes were provided to enable you to use channel exits written in languages other than Java:

- `MQExternalSecurityExit`, which implements the `MQSecurityExit` interface
- `MQExternalSendExit`, which implements the `MQSendExit` interface

- MQExternalReceiveExit, which implements the MQReceiveExit interface

The use of these classes remains valid but the new method is preferred.

To use a security exit that is not written in Java, an application first had to create an MQExternalSecurityExit object. The application specified, as parameters on the MQExternalSecurityExit constructor, the name of the library containing the security exit, the name of the entry point for the security exit, and the user data to be passed to the security exit when it was called. Channel exit programs that are not written in Java were stored in the directory shown in Table 153 on page 1065.

*Using a sequence of channel send or receive exits in IBM MQ classes for Java:*

An IBM MQ classes for Java application can use a sequence of channel send or receive exits that are run in succession.

To use a sequence of send exits, an application can create either a List or a String containing the send exits. If a List is used, each element of the List can be any of the following:

- An instance of a user defined class that implements the WMQSendExit interface
- An instance of a user defined class that implements the MQSendExit interface (for a send exit written in Java )
- An instance of the MQExternalSendExit class (for a send exit not written in Java )
- An instance of the MQSendExitChain class
- An instance of the String class

A List cannot contain another List.

The application can use a sequence of receive exits in a similar manner.

If a String is used, it must consist of one or more comma-separated exit definitions, each of which can be the name of a Java class, or a C program in the format `library(function)`.

The application then assigns the List or String object to the MQEnvironment.channelSendExit field before creating an MQQueueManager object.

The context of information passed to exits is solely within the domain of the exits. For example, if a Java exit and a C exit are chained, the presence of the Java exit has no effect on the C exit.

### **Using exit chain classes**

In versions of IBM MQ earlier than Version 7.0, two classes were provided to allow sequences of exits:

- MQSendExitChain, which implements the MQSendExit interface
- MQReceiveExitChain, which implements the MQReceiveExit interface

The use of these classes remains valid but the new method is preferred. Using the IBM MQ Classes for Java interfaces means that your application still has a dependency on `com.ibm.mq.jar`. If the new set of interfaces in the `com.ibm.mq.exits` package are used there is no dependency on `com.ibm.mq.jar`.

To use a sequence of send exits, an application created a list of objects, where each object was one of the following:

- An instance of a user defined class that implements the MQSendExit interface (for a send exit written in Java )
- An instance of the MQExternalSendExit class (for a send exit not written in Java )
- An instance of the MQSendExitChain class

The application created an MQSendExitChain object by passing this list of objects as a parameter on the constructor. The application would then have assigned the MQSendExitChain object to the

MQEnvironment.sendExit field before creating an MQQueueManager object.

### Channel compression in IBM MQ classes for Java:

Compressing the data that flows on a channel can improve the performance of the channel and reduce network traffic. IBM MQ classes for Java use the compression function built into IBM MQ.

Using function supplied with IBM MQ, you can compress the data that flows on message channels and MQI channels and, on either type of channel, you can compress header data and message data independently of each other. By default, no data is compressed on a channel. For a full description of channel compression, including how it is implemented in IBM MQ, see Data compression (COMPMSG) and Header compression (COMPHDR).

An IBM MQ classes for Java application specifies the techniques that can be used for compressing header or message data on a client connection by creating a java.util.Collection object. Each compression technique is an Integer object in the collection, and the order in which the application adds the compression techniques to the collection is the order in which the compression techniques are negotiated with the queue manager when the client connection starts. The application can then assign the collection to the hdrCompList field, for header data, or the msgCompList field, for message data, in the MQEnvironment class. When the application is ready, it can start the client connection by creating an MQQueueManager object.

The following code fragments illustrate the approach described. The first code fragment shows you how to implement header data compression:

```
Collection headerComp = new Vector();
headerComp.add(new Integer(CMQXC.MQCOMPRESS_SYSTEM));
:
MQEnvironment.hdrCompList = headerComp;
:
MQQueueManager qMgr = new MQQueueManager(QM);
```

The second code fragment shows you how to implement message data compression:

```
Collection msgComp = new Vector();
msgComp.add(new Integer(CMQXC.MQCOMPRESS_RLE));
msgComp.add(new Integer(CMQXC.MQCOMPRESS_ZLIBHIGH));
:
MQEnvironment.msgCompList = msgComp;
:
MQQueueManager qMgr = new MQQueueManager(QM);
```

In the second example, the compression techniques are negotiated in the order RLE, then ZLIBHIGH, when the client connection starts. The compression technique that is selected cannot be changed during the lifetime of the MQQueueManager object.

The compression techniques for header and message data that are supported by both the client and the queue manager on a client connection are passed to a channel exit as collections in the hdrCompList and msgCompList fields of an MQChannelDefinition object. The actual techniques that are currently being used for compressing header and message data on a client connection are passed to a channel exit in the CurHdrCompression and CurMsgCompression fields of an MQChannelExit object.

If compression is used on a client connection, the data is compressed before any channel send exits are processed and extracted after any channel receive exits are processed. The data passed to send and receive exits is therefore in a compressed state.

For more information about specifying compression techniques, and about which compression techniques are available, see Class com.ibm.mq.MQEnvironment and Interface com.ibm.mq.MQC.

## Sharing a TCP/IP connection in IBM MQ classes for Java:

Multiple instances of an MQI channel can be made to share a single TCP/IP connection.

In IBM MQ classes for Java, you use the `MQEnvironment.sharingConversations` variable to control the number of conversations that can share a single TCP/IP connection.

The `SHARECNV` attribute is a best effort approach to connection sharing. Therefore when a `SHARECNV` value greater than 0 is used with the IBM MQ classes for Java it is not guaranteed that a new connection request will always share an already established connection.

## Connection pooling in IBM MQ classes for Java:

IBM MQ classes for Java allows spare connections to be pooled for reuse.

IBM MQ classes for Java provides additional support for applications that deal with multiple connections to IBM MQ queue managers. When a connection is no longer required, instead of destroying it, it can be pooled and later reused. This can provide a substantial performance enhancement for applications and middleware that connect serially to arbitrary queue managers.

IBM MQ provides a default connection pool. Applications can activate or deactivate this connection pool by registering and deregistering tokens through the `MQEnvironment` class. If the pool is active when IBM MQ classes for Java constructs an `MQueueManager` object, it searches this default pool and reuses any suitable connection. When an `MQueueManager.disconnect()` call occurs, the underlying connection is returned to the pool.

Alternatively, applications can construct an `MQueueSimpleConnectionManager` connection pool for a particular use. Then, the application can either specify that pool during construction of an `MQueueManager` object, or pass that pool to `MQEnvironment` for use as the default connection pool.

To prevent connections from using too much resource, you can limit the total number of connections that an `MQueueSimpleConnectionManager` object can handle, and you can limit the size of the connection pool. Setting limits is useful if there are conflicting demands for connections within a JVM.

By default, the `getMaxConnections()` method returns the value zero, which means that there is no limit to the number of connections that the `MQueueSimpleConnectionManager` object can handle. You can set a limit by using the `setMaxConnections()` method. If you set a limit and the limit is reached, a request for a further connection might cause an `MQueueException` to be thrown, with a reason code of `MQRC_MAX_CONNS_LIMIT_REACHED`.

*Controlling the default connection pool in IBM MQ classes for Java:*

This example shows how to use the default connection pool.

Consider the following example application, `MQApp1`:

```
import com.ibm.mq.*;
public class MQApp1
{
 public static void main(String[] args) throws MQueueException
 {
 for (int i=0; i<args.length; i++) {
 MQueueManager qmgr=new MQueueManager(args[i]);
 :
 : (do something with qmgr)
 :
 }
 }
}
```

```

 qmgr.disconnect();
 }
}

```

MQApp1 takes a list of local queue managers from the command line, connects to each in turn, and performs some operation. However, when the command line lists the same queue manager many times, it is more efficient to connect only once, and to reuse that connection many times.

IBM MQ classes for Java provides a default connection pool that you can use to do this. To enable the pool, use one of the `MQEnvironment.addConnectionPoolToken()` methods. To disable the pool, use `MQEnvironment.removeConnectionPoolToken()`.

The following example application, MQApp2, is functionally identical to MQApp1, but connects only once to each queue manager.

```

import com.ibm.mq.*;
public class MQApp2
{
 public static void main(String[] args) throws MQException
 {
 MQPoolToken token=MQEnvironment.addConnectionPoolToken();

 for (int i=0; i<args.length; i++) {
 MQQueueManager qmgr=new MQQueueManager(args[i]);
 :
 : (do something with qmgr)
 :
 qmgr.disconnect();
 }

 MQEnvironment.removeConnectionPoolToken(token);
 }
}

```

The first bold line activates the default connection pool by registering an `MQPoolToken` object with `MQEnvironment`.

The `MQQueueManager` constructor now searches this pool for an appropriate connection and only creates a connection to the queue manager if it cannot find an existing one. The `qmgr.disconnect()` call returns the connection to the pool for later reuse. These API calls are the same as the sample application MQApp1.

The second highlighted line deactivates the default connection pool, which destroys any queue manager connections stored in the pool. This is important because otherwise the application would terminate with a number of live queue manager connections in the pool. This situation could cause errors that would appear in the queue manager logs.

If an application uses a client channel definition table (CCDT) to connect to a queue manager, the `MQQueueManager` constructor first searches the table for a suitable client connection channel definition. If one is found, the constructor searches the default connection pool for a connection that can be used for the channel. If the constructor cannot find a suitable connection in the pool, it then searches the client channel definition table for the next suitable client connection channel definition, and proceeds as described previously. If the constructor completes its search of the client channel definition table and fails to find any suitable connection in the pool, the constructor starts a second search of the table. During this search, the constructor tries to create a new connection for each suitable client connection channel definition in turn, and uses the first connection that it manages to create.

The default connection pool stores a maximum of ten unused connections, and keeps unused connections active for a maximum of five minutes. The application can alter this (for details, see “Supplying a different connection pool in IBM MQ classes for Java” on page 1072 ).

Instead of using MQEnvironment to supply an MQPoolToken, the application can construct its own:

```
MQPoolToken token=new MQPoolToken();
MQEnvironment.addConnectionPoolToken(token);
```

Some applications or middleware vendors provide subclasses of MQPoolToken in order to pass information to a custom connection pool. They can be constructed and passed to addConnectionPoolToken() in this way so that extra information can be passed to the connection pool.

*The default connection pool and multiple components in IBM MQ classes for Java:*

This example shows how to add or remove MQPoolTokens from a static set of registered MQPoolToken objects.

MQEnvironment holds a static set of registered MQPoolToken objects. To add or remove MQPoolTokens from this set, use the following methods:

- MQEnvironment.addConnectionPoolToken()
- MQEnvironment.removeConnectionPoolToken()

An application might consist of many components that exist independently and perform work using a queue manager. In such an application, each component should add an MQPoolToken to the MQEnvironment set for its lifetime.

For example, the example application MQApp3 creates ten threads and starts each one. Each thread registers its own MQPoolToken, waits for a length of time, then connects to the queue manager. After the thread disconnects, it removes its own MQPoolToken.

The default connection pool remains active while there is at least one token in the set of MQPoolTokens, so it will remain active for the duration of this application. The application does not need to keep a master object in overall control of the threads.

```
import com.ibm.mq.*;
public class MQApp3
{
 public static void main(String[] args)
 {
 for (int i=0; i<10; i++) {
 MQApp3_Thread thread=new MQApp3_Thread(i*60000);
 thread.start();
 }
 }
}

class MQApp3_Thread extends Thread
{
 long time;

 public MQApp3_Thread(long time)
 {
 this.time=time;
 }

 public synchronized void run()
 {
 MQPoolToken token=MQEnvironment.addConnectionPoolToken();
 try {
 wait(time);
 MQQueueManager qmgr=new MQQueueManager("my.qmgr.1");
```

```

 :
 : (do something with qmgr)
 :
 qmgr.disconnect();
 }
 catch (MQException mqe) {System.err.println("Error occurred!");}
 catch (InterruptedException ie) {}

 MQEnvironment.removeConnectionPoolToken(token);
}
}

```

*Supplying a different connection pool in IBM MQ classes for Java:*

This example shows how to use the class **com.ibm.mq.MQSimpleConnectionManager** to supply a different connection pool.

This class provides basic facilities for connection pooling, and applications can use this class to customize the behavior of the pool.

Once it is instantiated, an **MQSimpleConnectionManager** can be specified on the **MQQueueManager** constructor. The **MQSimpleConnectionManager** then manages the connection that underlies the constructed **MQQueueManager**. If the **MQSimpleConnectionManager** contains a suitable pooled connection, that connection is reused and returned to the **MQSimpleConnectionManager** after an **MQQueueManager.disconnect()** call.

The following code fragment demonstrates this behavior:

```

MQSimpleConnectionManager myConnMan=new MQSimpleConnectionManager();
myConnMan.setActive(MQSimpleConnectionManager.MODE_ACTIVE);
MQQueueManager qmgr=new MQQueueManager("my.qmgr.1", myConnMan);
:
: (do something with qmgr)
:
qmgr.disconnect();

MQQueueManager qmgr2=new MQQueueManager("my.qmgr.1", myConnMan);
:
: (do something with qmgr2)
:
qmgr2.disconnect();
myConnMan.setActive(MQSimpleConnectionManager.MODE_INACTIVE);

```

The connection that is forged during the first **MQQueueManager** constructor is stored in **myConnMan** after the **qmgr.disconnect()** call. The connection is then reused during the second call to the **MQQueueManager** constructor.

The second line enables the **MQSimpleConnectionManager**. The last line disables **MQSimpleConnectionManager**, destroying any connections held in the pool. An **MQSimpleConnectionManager** is, by default, in **MODE\_AUTO**, which is described later in this section.

An **MQSimpleConnectionManager** allocates connections on a most-recently-used basis, and destroys connections on a least-recently-used basis. By default, a connection is destroyed if it has not been used for five minutes, or if there are more than ten unused connections in the pool. You can alter these values by calling **MQSimpleConnectionManager.setTimeout()**.

You can also set up an **MQSimpleConnectionManager** for use as the default connection pool, to be used when no Connection Manager is supplied on the **MQQueueManager** constructor.

The following application demonstrates this:



```

import com.ibm.mq.*;
public class MQApp4
{
 public static void main(String []args)
 {
 MQSimpleConnectionManager myConnMan=new MQSimpleConnectionManager();
 myConnMan.setActive(MQSimpleConnectionManager.MODE_AUTO);
 myConnMan.setTimeout(3600000);
 myConnMan.setMaxConnections(75);
 myConnMan.setMaxUnusedConnections(50);
 MQEnvironment.setDefaultConnectionManager(myConnMan);
 MQApp3.main(args);
 }
}

```

The bold lines create and configure an MQSimpleConnectionManager object. The configuration does the following:

- Ends connections that are not used for an hour
- Limits the number of connections managed by myConnMan to 75
- Limits the number of unused connections in the pool to 50
- Sets MODE\_AUTO, which is the default. This means that the pool is active only if it is the default connection manager, and there is at least one token in the set of MQPoolTokens held by MQEnvironment.

The new MQSimpleConnectionManager is then set as the default connection manager.

In the last line, the application calls MQApp3.main(). This runs a number of threads, where each thread uses IBM MQ independently. These threads use myConnMan when they forge connections.

### **JTA/JDBC coordination using IBM MQ classes for Java:**

IBM MQ classes for Java supports the MQQueueManager.begin() method, which allows IBM MQ to act as a coordinator for a database which provides a JDBC type 2 or JDBC type 4 compliant driver.

This support is not available on all platforms. To check which platforms support JDBC coordination, see System Requirements for IBM MQ.

To use the XA-JTA support, you must use the special JTA switch library. The method for using this library varies depending on whether you are using Windows or one of the other platforms.

*Configuring JTA/JDBC coordination on Windows:*

The XA library is supplied as a DLL with a name of the format jdbcxxx.dll.

The supplied jdbcora12.dll provides compatibility with Oracle 12C, for an IBM MQ Windows server installation.

On Windows systems, the XA library is supplied as a complete DLL. The name of this DLL is jdbcxxx.dll where xxx indicates the database for which the switch library has been compiled. This library is in the java\lib\jdbc or java\lib64\jdbc directory of your IBM MQ classes for Java installation. You must declare the XA library, also described as the switch load file, to the queue manager. Use the IBM MQ Explorer. Specify the details of the switch load file in the queue manager properties panel, under XA resource manager. You must only give the name of the library. For example:

For a Db2 database set the SwitchFile field to: dbcdb2

For an Oracle database set the SwitchFile field to: jdbcora

*Configuring JTA/JDBC coordination on platforms other than Windows:*

Object files are supplied. Link the appropriate one using the supplied makefile, and declare it to the queue manager using the configuration file.

For each database management system, IBM MQ provides two object files. You must link one object file to create a 32-bit switch library, and link the other object file to create a 64-bit switch library. For Db2, the name of each object file is `jdbcdb2.o` and, for Oracle, the name of each object file is `jdbcora.o`.

You must link each object file using the appropriate makefile supplied with IBM MQ. A switch library requires other libraries, which might be stored in different locations on different systems. However, a switch library cannot use the library path environment variable to locate these libraries because the switch library is loaded by the queue manager, which runs in a `setuid` environment. The supplied makefile therefore ensures that a switch library contains the fully qualified path names of these libraries.

To create a switch library, enter a **make** command with the following format. To create a 32-bit switch library, enter the command in the `/java/lib/jdbc` directory of your IBM MQ installation. To create a 64-bit switch library, enter the command in the `/java/lib64/jdbc` directory.

```
make DBMS
```

where *DBMS* is the database management system for which you are creating the switch library. The valid values are `db2` for Db2 and `oracle` for Oracle.

Here is an example of a **make** command:

```
make db2
```

Note the following points:

- To run 32-bit applications, you must create both a 32-bit and a 64-bit switch library for each database management system that you are using. To run 64-bit applications, you need create only a 64-bit switch library. For Db2, the name of each switch library is `jdbcdb2` and, for Oracle, the name of each switch library is `jdbcora`. The makefiles ensure that 32-bit and 64-bit switch libraries are stored in different IBM MQ directories. A 32-bit switch library is stored in the `/java/lib/jdbc` directory, and a 64-bit switch library is stored in the `/java/lib64/jdbc` directory.
- Because you can install Oracle anywhere on a system, the makefiles use the `ORACLE_HOME` environment variable to locate where Oracle is installed.

After you have created the switch libraries for Db2, Oracle, or both, you must declare them to your queue manager. If the queue manager configuration file (`qm.ini`) already contains `XAResourceManager` stanzas for Db2 or Oracle databases, you must replace the `SwitchFile` entry in each stanza by one of the following:

**For a Db2 database**

```
SwitchFile=jdbcdb2
```

**For an Oracle database**

```
SwitchFile=jdbcora
```

Do not specify the fully qualified path name of either the 32-bit or 64-bit switch library. Specify only the name of the library.

If the queue manager configuration file does not already contain `XAResourceManager` stanzas for Db2 or Oracle databases, or if you want to add additional `XAResourceManager` stanzas, see *Administering* for information about how to construct an `XAResourceManager` stanza. However, each `SwitchFile` entry in a new `XAResourceManager` stanza must be exactly as described previously for a Db2 or Oracle database. You must also include the entry `ThreadOfControl=PROCESS`.

After you have updated the queue manager configuration file, and made sure that all appropriate database environment variables have been set, you can restart the queue manager.

*Using JTA/JDBC coordination:*

Code your API calls as in the supplied example.

The basic sequence of API calls for a user application is:

```
qMgr = new MQQueueManager("QM1")
Connection con = qMgr.getJDBCConnection(xads);
qMgr.begin()
```

< Perform MQ and DB operations to be grouped in a unit of work >

```
qMgr.commit() or qMgr.backout();
con.close()
qMgr.disconnect()
```

xads in the getJDBCConnection call is a database-specific implementation of the XADataSource interface, which defines the details of the database to connect to. See the documentation for your database to determine how to create an appropriate XADataSource object to pass into getJDBCConnection.

You must also update your class path with the appropriate database-specific jar files for performing JDBC work.

If you must connect to multiple databases, you must call getJDBCConnection several times to perform the transaction across several different connections.

There are two forms of the getJDBCConnection, reflecting the two forms of XADataSource.getXAConnection:

```
public java.sql.Connection getJDBCConnection(javax.sql.XADataSource xads)
 throws MQException, SQLException, Exception

public java.sql.Connection getJDBCConnection(XADataSource dataSource,
 String userid, String password)
 throws MQException, SQLException, Exception
```

These methods declare Exception in their throws clauses to avoid problems with the JVM verifier for customers who are not using the JTA functions. The actual exception thrown is javax.transaction.xa.XAException which requires the jta.jar file to be added to the class path for programs that did not previously require it.

To use the JTA/JDBC support, you must include the following statement in your application:

```
MQEnvironment.properties.put(CMQC.THREAD_AFFINITY_PROPERTY, new Boolean(true));
```

*Known problems and limitations with JTA/JDBC coordination:*

Some of the problems and limitations of JTA/JDBC support depend on the database management system in use, for example, tested JDBC drivers behave differently when the database is shut down while an application is running. If the connection to the database that an application is using is broken, there are steps that the application can perform to reestablish a new connection to the queue manager and the database so that it can then use those new connections to perform the transactional work required.

Because the JTA/JDBC support makes calls to JDBC drivers, the implementation of those JDBC drivers can have a significant effect on the system behavior. In particular, tested JDBC drivers behave differently when the database is shut down while an application is running.

**Important:** Always avoid abruptly shutting down a database while there are applications that are holding open connections to it.

**Note:** An IBM MQ classes for Java application must connect by using bindings mode to make IBM MQ act as a database coordinator.

### **Multiple XAResourceManager stanzas**

The use of more than one XAResourceManager stanza in a queue manager configuration file, `qm.ini`, is not supported. Any XAResourceManager stanza other than the first is ignored.

### **Db2**

Sometimes Db2 returns a SQL0805N error. This problem can be resolved with the following CLP command:

```
DB2 bind @db2cli.lst blocking all grant public
```

For more information, refer to the Db2 documentation.

The XAResourceManager stanza must be configured to use `ThreadOfControl=PROCESS`. For Db2 Version 8.1 and higher, this does not match the default thread of control setting for Db2, so `toc=p` must be specified in the XA Open String. An example XAResourceManager stanza for Db2 with JTA/JDBC coordination is as follows:

```
XAResourceManager:
 Name=jdbcdb2
 SwitchFile=jdbcdb2
 XAOpenString=uid=userid,db=dbalias,pwd=password,toc=p
 ThreadOfControl=PROCESS
```

This does not prevent the Java applications that use JTA/JDBC coordination from being multithreaded themselves.

**Oracle** Calling the `JDBC Connection.close()` method after `MQQueueManager.disconnect()` generates an `SQLException`. Either call `Connection.close()` before `MQQueueManager.disconnect()`, or omit the call to `Connection.close()`.

### **Handling issues with database connections**

When an IBM MQ classes for Java application uses the JTA/JDBC support that is provided by IBM MQ, it typically performs the following steps:

1. Creates a new `MQQueueManager` object to represent a connection to the queue manager that will act as the transaction manager.
2. Constructs an `XADataSource` object that contains details about how to connect to the database that will be enlisted in the transaction.
3. Calls the method `MQQueueManager.getJDBCConnection(XADataSource)` passing in the `XADataSource` that was created previously. This causes the IBM MQ classes for Java to establish a connection to the database.

4. Calls the method `MQQueueManager.begin()` to start the XA transaction.
5. Performs the messaging and database work.
6. When all of the required work has been completed, calls the method `MQQueueManager.commit()`. This completes the XA transaction.
7. If a new XA transaction is required at this point, the application can repeat steps 4, 5 and 6.
8. When the application has finished, it should close the database connection that was created at step 3, and then call the method `MQQueueManager.disconnect()` to disconnect from the queue manager.

The IBM MQ classes for Java maintain an internal list of all of the database connections that have been created when an application calls `MQQueueManager.getJDBCConnection(XADataSource)`. If a queue manager needs to communicate with the database during the processing of the XA transaction, the following processing takes place:

1. The queue manager calls into the IBM MQ classes for Java, passing in details of the XA call that needs to be passed to the database.
2. The IBM MQ classes for Java then look up the appropriate connection in the list, and then use that connection to flow the XA call to the database.

If the connection to the database is lost at any point during this processing, the application should:

1. Back out any existing work that was done under the transaction, by calling the method `MQQueueManager.backout()`.
2. Close the database connection. This should cause the IBM MQ classes for Java to remove details of the broken database connection from its internal list.
3. Disconnect from the queue manager, by calling the method `MQQueueManager.disconnect()`.
4. Establish a new connection to the queue manager, by constructing a new `MQQueueManager` object.
5. Create a new database connection, by calling the method `MQQueueManager.getJDBCConnection(XADataSource)`.
6. Perform the transactional work again.

This allows the application to reestablish a new connection to the queue manager and the database, and then use those connections to perform the transactional work required.

### **Transport Layer Security (TLS) support in IBM MQ classes for Java:**

IBM MQ classes for Java client applications support TLS encryption. You require a JSSE provider to use TLS encryption.

IBM MQ classes for Java client applications using `TRANSPORT(CLIENT)` support TLS encryption. TLS provides communication encryption, authentication, and message integrity. It is typically used to secure communications between any two peers on the Internet or within an intranet.

IBM MQ classes for Java uses Java Secure Socket Extension (JSSE) to handle TLS encryption, and so requires a JSSE provider. JSE v1.4 JVMs have a JSSE provider built in. Details of how to manage and store certificates can vary from provider to provider. For information about this, refer to your JSSE provider's documentation.

This section assumes that your JSSE provider is correctly installed and configured, and that suitable certificates have been installed and made available to your JSSE provider.

If your IBM MQ classes for Java client application uses a client channel definition table (CCDT) to connect to a queue manager, see “Using a client channel definition table with IBM MQ classes for Java” on page 1047.

### *Enabling TLS in IBM MQ classes for Java:*

To enable TLS, you specify a CipherSuite. There are two ways of specifying a CipherSuite.

TLS is supported only for client connections. To enable TLS, you must specify the CipherSuite to use when communicating with the queue manager, and this CipherSuite must match the CipherSpec set on the target channel. Additionally, the named CipherSuite must be supported by your JSSE provider. However, CipherSuites are distinct from CipherSpecs and so have different names. “TLS CipherSpecs and CipherSuites in IBM MQ classes for Java” on page 1083 contains a table mapping the CipherSpecs supported by IBM MQ to their equivalent CipherSuites as known to JSSE.

To enable TLS, specify the CipherSuite using the `sslCipherSuite` static member variable of `MQEnvironment`. The following example attaches to an `SVRCONN` channel named `SECURE.SVRCONN.CHANNEL`, which has been set up to require TLS with a CipherSpec of `TLS_RSA_WITH_AES_128_CBC_SHA`:

```
MQEnvironment.hostname = "your_hostname";
MQEnvironment.channel = "SECURE.SVRCONN.CHANNEL";
MQEnvironment.sslCipherSuite = "SSL_RSA_WITH_AES_128_CBC_SHA";
MQQueueManager qmgr = new MQQueueManager("your_Q_manager");
```

Although the channel has a CipherSpec of `TLS_RSA_WITH_AES_128_CBC_SHA`, the Java application must specify a CipherSuite of `SSL_RSA_WITH_AES_128_CBC_SHA`. See “TLS CipherSpecs and CipherSuites in IBM MQ classes for Java” on page 1083 for a list of mappings between CipherSpecs and CipherSuites.

An application can also specify a CipherSuite by setting the environment property `CMQC.SSL_CIPHER_SUITE_PROPERTY`.

Alternatively, use the Client Channel Definition Table (CCDT). For more information, see “Using a client channel definition table with IBM MQ classes for Java” on page 1047

If you require a client connection to use a CipherSuite that is supported by the IBM Java JSSE FIPS provider (IBMJSSEFIPS), an application can set the `sslFipsRequired` field in the `MQEnvironment` class to `true`. Alternatively, the application can set the environment property `CMQC.SSL_FIPS_REQUIRED_PROPERTY`. The default value is `false`, which means that a client connection can use any CipherSuite that is supported by IBM MQ.

If an application uses more than one client connection, the value of the `sslFipsRequired` field that is used when the application creates the first client connection determines the value that is used when the application creates any subsequent client connection. Therefore when the application creates a subsequent client connection, the value of the `sslFipsRequired` field is ignored. You must restart the application if you want to use a different value for the `sslFipsRequired` field.

To connect successfully using TLS, the JSSE truststore must be set up with certificate authority root certificates from which the certificate presented by the queue manager can be authenticated. Similarly, if `SSLClientAuth` on the `SVRCONN` channel has been set to `MQSSL_CLIENT_AUTH_REQUIRED`, the JSSE keystore must contain an identifying certificate that is trusted by the queue manager.

## Related information:

Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows

*Using the distinguished name of the queue manager in IBM MQ classes for Java:*

The queue manager identifies itself using a TLS certificate, which contains a distinguished name (DN). An IBM MQ classes for Java client application can use this DN to ensure that it is communicating with the correct queue manager.

A DN pattern is specified using the `sslPeerName` variable of `MQEnvironment`. For example, setting:  
`MQEnvironment.sslPeerName = "CN=QMGR.*, OU=IBM, OU=WEBSPHERE";`

allows the connection to succeed only if the queue manager presents a certificate with a Common Name beginning `QMGR.`, and at least two Organizational Unit names, the first of which must be `IBM` and the second `WebSphere`.

If `sslPeerName` is set, connections succeed only if it is set to a valid pattern and the queue manager presents a matching certificate.

An application can also specify the distinguished name of the queue manager by setting the environment property `CMQC.SSL_PEER_NAME_PROPERTY`. For more information about distinguished names, see [Distinguished names](#).

*Using certificate revocation lists in IBM MQ classes for Java:*

Specify the certificate revocation lists to use through the `java.security.cert.CertStore` class. IBM MQ classes for Java then checks certificates against the specified CRL.

A certificate revocation list (CRL) is a set of certificates that have been revoked, either by the issuing certificate authority or by the local organization. CRLs are typically hosted on LDAP servers. With Java 2 v1.4, a CRL server can be specified at connect-time and the certificate presented by the queue manager is checked against the CRL before the connection is allowed. For more information about certificate revocation lists and IBM MQ, see [Working with Certificate Revocation Lists and Authority Revocation Lists and Accessing CRLs and ARLs with IBM MQ classes for Java and IBM MQ classes for JMS](#).

**Note:** To use a `CertStore` successfully with a CRL hosted on an LDAP server, make sure that your Java Software Development Kit (SDK) is compatible with the CRL. Some SDKs require that the CRL conforms to RFC 2587, which defines a schema for LDAP v2. Most LDAP v3 servers use RFC 2256 instead.

The CRLs to use are specified through the `java.security.cert.CertStore` class. Refer to documentation on this class for full details of how to obtain instances of `CertStore`. To create a `CertStore` based on an LDAP server, first create an `LDAPCertStoreParameters` instance, initialized with the server and port settings to use. For example:

```
import java.security.cert.*;
CertStoreParameters csp = new LDAPCertStoreParameters("cr1_server", 389);
```

Having created a `CertStoreParameters` instance, use the static constructor on `CertStore` to create a `CertStore` of type LDAP:

```
CertStore cs = CertStore.getInstance("LDAP", csp);
```

Other `CertStore` types (for example, `Collection`) are also supported. Commonly there are several CRL servers set up with identical CRL information to give redundancy. When you have a `CertStore` object for each of these CRL servers, place them all in a suitable `Collection`. The following example shows the `CertStore` objects placed in an `ArrayList`:

```
import java.util.ArrayList;
Collection crls = new ArrayList();
crls.add(cs);
```

This Collection can be set into the MQEnvironment static variable, sslCertStores, before connecting to enable CRL checking:

```
MQEnvironment.sslCertStores = crls;
```

The certificate presented by the queue manager when a connection is being set up is validated as follows:

1. The first CertStore object in the Collection identified by sslCertStores is used to identify a CRL server.
2. An attempt is made to contact the CRL server.
3. If the attempt is successful, the server is searched for a match for the certificate.
  - a. If the certificate is found to be revoked, the search process is over and the connection request fails with reason code MQRC\_SSL\_CERTIFICATE\_REVOKED.
  - b. If the certificate is not found, the search process is over and the connection is allowed to proceed.
4. If the attempt to contact the server is unsuccessful, the next CertStore object is used to identify a CRL server and the process repeats from step 2.

If this was the last CertStore in the Collection, or if the Collection contains no CertStore objects, the search process failed, and the connection request fails with reason code MQRC\_SSL\_CERT\_STORE\_ERROR.

The Collection object determines the order in which CertStores are used.

The Collection of CertStores can also be set using the CMQC.SSL\_CERT\_STORE\_PROPERTY. As a convenience, this property also allows a single CertStore to be specified without being a member of a Collection.

If sslCertStores is set to null, no CRL checking is performed. This property is ignored if sslCipherSuite is not set.

*Renegotiating the secret key in IBM MQ classes for Java:*

An IBM MQ classes for Java client application can control when the secret key that is used for encryption on a client connection is renegotiated, in terms of the total number of bytes sent and received.

The application can do this in either of the following ways: If the application uses more than one of these ways, the usual precedence rules apply.

- By setting the sslResetCount field in the MQEnvironment class.
- By setting the environment property MQC.SSL\_RESET\_COUNT\_PROPERTY in a Hashtable object. The application then assigns the hashtable to the properties field in the MQEnvironment class, or passes the hashtable to an MQQueueManager object on its constructor.

The value of the sslResetCount field or environment property MQC.SSL\_RESET\_COUNT\_PROPERTY represents the total number of bytes sent and received by the IBM MQ classes for Java client code before the secret key is renegotiated. The number of bytes sent is the number before encryption, and the number of bytes received is the number after decryption. The number of bytes also includes control information sent and received by the IBM MQ classes for Java client.

If the reset count is zero, which is the default value, the secret key is never renegotiated. The reset count is ignored if no CipherSuite is specified.



*Supplying a customized SSLSocketFactory in IBM MQ classes for Java:*

If you use a customized JSSE Socket Factory, set the `MQEnvironment.sslSocketFactory` to the customized factory object. Details vary between different JSSE implementations.

Different JSSE implementations can provide different features. For example, a specialized JSSE implementation might allow configuration of a particular model of encryption hardware. Additionally, some JSSE providers allow customization of keystores and truststores by program, or allow the choice of identity certificate from the keystore to be altered. In JSSE, all these customizations are abstracted into a factory class, `javax.net.ssl.SSLSocketFactory`.

See your JSSE documentation for details of how to create a customized `SSLSocketFactory` implementation. The details vary from provider to provider, but a typical sequence of steps might be:

1. Create an `SSLContext` object using a static method on `SSLContext`
2. Initialize this `SSLContext` with appropriate `KeyManager` and `TrustManager` implementations (created from their own factory classes)
3. Create an `SSLSocketFactory` from the `SSLContext`

When you have an `SSLSocketFactory` object, set the `MQEnvironment.sslSocketFactory` to the customized factory object. For example:

```
javax.net.ssl.SSLSocketFactory sf = sslContext.getSocketFactory();
MQEnvironment.sslSocketFactory = sf;
```

IBM MQ classes for Java use this `SSLSocketFactory` to connect to the IBM MQ queue manager. This property can also be set using the `CMQC.SSL_SOCKET_FACTORY_PROPERTY`. If `sslSocketFactory` is set to null, the default `SSLSocketFactory` of the JVM is used. This property is ignored if `sslCipherSuite` is not set.

When you use custom `SSLSocketFactories`, consider the effect of TCP/IP connection sharing. If connection sharing is possible then a new socket is not requested of the `SSLSocketFactory` supplied, even if the socket produced would be different in some way in the context of a subsequent connection request. For example, if a different client certificate is to be presented on a subsequent connection, then connection sharing must not be allowed.

*Making changes to the JSSE keystore or truststore in IBM MQ classes for Java:*

If you change the JSSE keystore or truststore, you must perform certain actions for the changes to take effect.

If you change the contents of the JSSE keystore or truststore, or change the location of the keystore or truststore file, IBM MQ classes for Java applications that are running at the time do not automatically pick up the changes. For the changes to take effect, the following actions must be performed:

- The applications must close all their connections, and destroy any unused connections in connection pools.
- If your JSSE provider caches information from the keystore and truststore, this information must be refreshed.

After these actions have been performed, the applications can then re-create their connections.

Depending on how you design your applications, and on the function provided by your JSSE provider, it might be possible to perform these actions without stopping and restarting your applications. However, stopping and restarting the applications might be the simplest solution.

*Error handling when using TLS with IBM MQ classes for Java:*

A number of reason codes can be issued by IBM MQ classes for Java when connecting to a queue manager using TLS.

These are explained in the following list:

**MQRC\_SSL\_NOT\_ALLOWED**

The `sslCipherSuite` property was set, but `bindings.connect` was used. Only `client.connect` supports TLS.

**MQRC\_JSSE\_ERROR**

The JSSE provider reported an error that could not be handled by IBM MQ. This could be caused by a configuration problem with JSSE, or because the certificate presented by the queue manager could not be validated. The exception produced by JSSE can be retrieved using the `getCause()` method on `MQException`.

**MQRC\_SSL\_INITIALIZATION\_ERROR**

An `MQCONN` or `MQCONNX` call was issued with TLS configuration options specified, but an error occurred during the initialization of the TLS environment.

**MQRC\_SSL\_PEER\_NAME\_MISMATCH**

The DN pattern specified in the `sslPeerName` property did not match the DN presented by the queue manager.

**MQRC\_SSL\_PEER\_NAME\_ERROR**

The DN pattern specified in the `sslPeerName` property was not valid.

**MQRC\_UNSUPPORTED\_CIPHER\_SUITE**

The `CipherSuite` named in `sslCipherSuite` was not recognized by the JSSE provider. A full list of `CipherSuites` supported by the JSSE provider can be obtained by a program using the `SSLContextFactory.getSupportedCipherSuites()` method. A list of `CipherSuites` that can be used to communicate with IBM MQ can be found in "TLS CipherSpecs and CipherSuites in IBM MQ classes for Java" on page 1083.

**MQRC\_SSL\_CERTIFICATE\_REVOKED**

The certificate presented by the queue manager was found in a CRL specified with the `sslCertStores` property. Update the queue manager to use trusted certificates.

**MQRC\_SSL\_CERT\_STORE\_ERROR**

None of the supplied `CertStores` could be searched for the certificate presented by the queue manager. The `MQException.getCause()` method returns the error that occurred while searching the first `CertStore` attempted. If the causal exception is `NoSuchElementException`, `ClassCastException`, or `NullPointerException`, check that the `Collection` specified on the `sslCertStores` property contains at least one valid `CertStore` object.

*TLS CipherSpecs and CipherSuites in IBM MQ classes for Java:*

The ability of IBM MQ classes for Java applications to establish connections to a queue manager, depends on the CipherSpec specified at the server end of the MQI channel and the CipherSuite specified at the client end.

The following table lists the CipherSpecs supported by IBM MQ and their equivalent CipherSuites.

You should review the topic *Deprecated CipherSpecs* to see if any of the CipherSpecs, listed in the following table, have been deprecated by IBM MQ and, if so, at which update the CipherSpec was deprecated.

**Important:** The CipherSuites listed are those supported by the IBM Java Runtime Environment (JRE) supplied with IBM MQ. The CipherSuites that are listed include those supported by the Oracle Java JRE. For more information about configuring your application to use an Oracle Java JRE, see “Configuring your application to use IBM Java or Oracle Java CipherSuite mappings” on page 1085.

The table also indicates the protocol that is used for the communication, and whether or not the CipherSuite conforms to the FIPS 140-2 standard.

Ciphersuites denoted as FIPS 140-2 compliant can be used if the application has not been configured to enforce FIPS 140-2 compliance, but if FIPS 140-2 compliance has been configured for the application (see the following notes on configuration) only those CipherSuites which are marked as FIPS 140-2 compatible can be configured; attempting to use other CipherSuites results in an error.

**Note:** Each JRE can have multiple cryptographic security providers, each of which can contribute an implementation of the same CipherSuite. However, not all security providers are FIPS 140-2 certified. If FIPS 140-2 compliance is not enforced for an application then it is possible that an uncertified implementation of the CipherSuite might be used. Uncertified implementations might not operate in compliance with FIPS 140-2, even if the CipherSuite theoretically meets the minimum security level required by the standard. See the following notes for more information about configuring FIPS 140-2 enforcement in IBM MQ Java applications.

For more information about FIPS 140-2 and Suite-B compliance for CipherSpecs and CipherSuites, see *Specifying CipherSpecs*. You might also need to be aware of information that concerns US Federal Information Processing Standards.

To use the full set of CipherSuites and to operate with certified FIPS 140-2 and/or Suite-B compliance, a suitable JRE is required. IBM Java 7 Service Refresh 4 Fix Pack 2 or a higher level of IBM JRE provides the appropriate support.

**Note:** To use some CipherSuites, the 'unrestricted' policy files need to be configured in the JRE. For more details of how policy files are set up in an SDK or JRE, see the *IBM SDK Policy files* topic in the *Security Reference for IBM SDK, Java Technology Edition* for the version you are using..

Table 154. CipherSpecs supported by IBM MQ and their equivalent CipherSuites

CipherSpec	Equivalent CipherSuite (IBM JRE)	Equivalent CipherSuite (Oracle JRE)	Protocol	FIPS 140-2 compatible
ECDHE_ECDSA_3DES_EDE_CBC_SHA256	SSL_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA256	TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA256	TLS	Yes
ECDHE_ECDSA_AES_128_CBC_SHA256	SSL_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLS	Yes
ECDHE_ECDSA_AES_128_GCM_SHA256	SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLS	Yes
ECDHE_ECDSA_AES_256_CBC_SHA384	SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLS	Yes
ECDHE_ECDSA_AES_256_GCM_SHA384	SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLS	Yes
ECDHE_ECDSA_NULL_SHA256	SSL_ECDHE_ECDSA_WITH_NULL_SHA	TLS_ECDHE_ECDSA_WITH_NULL_SHA	TLS	Yes
ECDHE_ECDSA_RC4_128_SHA256	SSL_ECDHE_ECDSA_WITH_RC4_128_SHA	TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	TLS	Yes
ECDHE_RSA_3DES_EDE_CBC_SHA256	SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA256	TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA256	TLS	Yes
ECDHE_RSA_AES_128_CBC_SHA256	SSL_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLS	Yes
ECDHE_RSA_AES_128_GCM_SHA256	SSL_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLS	Yes
ECDHE_RSA_AES_256_CBC_SHA384	SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLS	Yes
ECDHE_RSA_AES_256_GCM_SHA384	SSL_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLS	Yes
ECDHE_RSA_NULL_SHA256	SSL_ECDHE_RSA_WITH_NULL_SHA	TLS_ECDHE_RSA_WITH_NULL_SHA	TLS	Yes
ECDHE_RSA_RC4_128_SHA256	SSL_ECDHE_RSA_WITH_RC4_128_SHA	TLS_ECDHE_RSA_WITH_RC4_128_SHA	TLS	Yes
TLS_RSA_WITH_3DES_EDE_CBC_SHA1	SSL_RSA_WITH_3DES_EDE_CBC_SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	TLS	Yes
TLS_RSA_WITH_AES_128_CBC_SHA	SSL_RSA_WITH_AES_128_CBC_SHA	TLS_RSA_WITH_AES_128_CBC_SHA	TLS	Yes
TLS_RSA_WITH_AES_128_CBC_SHA256	SSL_RSA_WITH_AES_128_CBC_SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256	TLS	Yes
TLS_RSA_WITH_AES_128_GCM_SHA256	SSL_RSA_WITH_AES_128_GCM_SHA256	TLS_RSA_WITH_AES_128_GCM_SHA256	TLS	Yes
TLS_RSA_WITH_AES_256_CBC_SHA	SSL_RSA_WITH_AES_256_CBC_SHA	TLS_RSA_WITH_AES_256_CBC_SHA	TLS	Yes
TLS_RSA_WITH_AES_256_CBC_SHA256	SSL_RSA_WITH_AES_256_CBC_SHA256	TLS_RSA_WITH_AES_256_CBC_SHA256	TLS	Yes
TLS_RSA_WITH_AES_256_GCM_SHA384	SSL_RSA_WITH_AES_256_GCM_SHA384	TLS_RSA_WITH_AES_256_GCM_SHA384	TLS	Yes
TLS_RSA_WITH_DES_CBC_SHA	SSL_RSA_WITH_DES_CBC_SHA	SSL_RSA_WITH_DES_CBC_SHA	SSL	No
TLS_RSA_WITH_NULL_SHA256	SSL_RSA_WITH_NULL_SHA256	TLS_RSA_WITH_NULL_SHA256	TLS	Yes
TLS_RSA_WITH_RC4_128_SHA	SSL_RSA_WITH_RC4_128_SHA	SSL_RSA_WITH_RC4_128_SHA	SSL	No

**Notes:**

1. This CipherSpec TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA is deprecated. However, it can still be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, you need to either avoid using triple DES, or enable secret key reset when using this CipherSpec.

**Configuring Ciphersuites and FIPS-compliance in an IBM MQ classes for Java application**

- An application that uses IBM MQ classes for Java can use either of two methods to set the CipherSuite for a connection:
  - Set the sslCipherSuite field in the MQEnvironment class to the CipherSuite name.
  - Set the property CMQC.SSL\_CIPHER\_SUITE\_PROPERTY in the properties hashtable passed to the MQQueueManager constructor to the CipherSuite name.
- An application that uses IBM MQ classes for Java can use either of two methods to enforce FIPS 140-2 compliance:

- Set the `sslFipsRequired` field to `true` in the `MQEnvironment` class.
- Set the property `CMQC.SSL_FIPS_REQUIRED_PROPERTY` in the properties hash table passed to the `MQQueueManager` constructor to `true`.

### Configuring your application to use IBM Java or Oracle Java CipherSuite mappings

You can configure whether your application uses the default IBM Java CipherSuite to IBM MQ CipherSpec mappings, or the Oracle CipherSuite to IBM MQ CipherSpec mappings. Therefore, you can use TLS CipherSuites whether your application uses an IBM JRE or an Oracle JRE. The Java System Property `com.ibm.mq.cfg.useIBMCipherMappings` controls which mappings are used. The property can be one of the following values:

**true** Use the IBM Java CipherSuite to IBM MQ CipherSpec mappings.

This value is the default value.

**false** Use the Oracle CipherSuite to IBM MQ CipherSpec mappings.

For more information about using IBM MQ Java and TLS Ciphers, see the MQdev blog posts [MQ Java](#), [TLS Ciphers](#), [Non-IBM JREs & APARs IT06775](#), [IV66840](#), [IT09423](#), [IT10837](#), and [The relationship between MQ CipherSpecs and Java Cipher Suites](#).

### Interoperability limitations

Certain CipherSuites might be compatible with more than one IBM MQ CipherSpec, depending on the protocol in use. However, only the CipherSuite/CipherSpec combination that uses the TLS version specified in Table 1 is supported. Attempting to use the unsupported combinations of CipherSuites and CipherSpecs will fail with an appropriate exception. Installations using any of these CipherSuite/CipherSpec combinations should move to a supported combination.

The following table shows the CipherSuites to which this limitation applies.

*Table 155. CipherSuites and their supported and unsupported CipherSpecs*

CipherSuite	Supported TLS CipherSpec	Unsupported SSL CipherSpec
SSL_RSA_WITH_3DES_EDE_CBC_SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA <sup>1</sup>	TRIPLE_DES_SHA_US
SSL_RSA_WITH_DES_CBC_SHA	TLS_RSA_WITH_DES_CBC_SHA	DES_SHA_EXPORT
SSL_RSA_WITH_RC4_128_SHA	TLS_RSA_WITH_RC4_128_SHA256	RC4_SHA_US

#### Note:

1. This CipherSpec `TLS_RSA_WITH_3DES_EDE_CBC_SHA` is deprecated. However, it can still be used to transfer up to 32 GB of data before the connection is terminated with error `AMQ9288`. To avoid this error, you need to either avoid using triple DES, or enable secret key reset when using this CipherSpec.

## Running IBM MQ classes for Java applications:

If you write an application (a class that contains a `main()` method), using either the client or the bindings mode, run your program using the Java interpreter.

Use the command:

```
java -Djava.library.path= library_path MyClass
```

where *library\_path* is the path to the IBM MQ classes for Java libraries. For more information, see “IBM MQ classes for Java libraries” on page 1030.

### Related information:

Tracing IBM MQ classes for Java applications

Tracing the IBM MQ Resource Adapter

## Java and JMS client connectivity to batch applications running on z/OS:

z/OS

V 9.0.4

### MQ Adv. VUE

A JMS, or IBM MQ classes for Java, application on z/OS can connect to a queue manager on z/OS, that has the **ADVCAP(ENABLED)** attribute, by using a client connection.

A value of **ADVCAP(ENABLED)** applies only to a z/OS queue manager, licensed as IBM MQ Advanced for z/OS, Value Unit Edition (see IBM MQ for z/OS product identifiers), running with **QMGRPROD** set to **ADVANCEDVUE**.

See **DISPLAY QMGR** for more information on **ADVCAP** and **START QMGR** for more information on **QMGRPROD**.

Note that batch is the only environment supported; there is no support for JMS for CICS or JMS for IMS.

An IBM MQ classes for JMS, or IBM MQ classes for Java, application on z/OS cannot connect, using client mode connection, to a queue manager that is not running on z/OS, or to a queue manager that does not have the **ADVCAP(ENABLED)** option.

If a JMS application attempts to connect using client mode, and the application is not allowed to do so, exception message **JMSFMQ0005** is issued.

If an IBM MQ classes for Java application on z/OS attempts to connect using client mode, and it is not allowed to do so, **MQRC\_ENVIRONMENT\_ERROR** is returned.

V 9.0.5

## Advanced Message Security (AMS) support

From IBM MQ Version 9.0.5, IBM MQ classes for JMS or IBM MQ classes for Java client applications can use AMS when connecting to IBM MQ Advanced for z/OS, Value Unit Edition queue managers on remote z/OS systems.

A new key store type, `jceracfks`, is supported in `keystore.conf` on z/OS only, where:

- The property name prefix is `jceracfks` and this name prefix is case insensitive.
- The key store is a RACF keyring.
- Passwords are not required, and will be ignored. This is because RACF keyrings do not use passwords.
- If you specify the provider, the provider must be `IBMJCE`.

When you use jceracfs with AMS, the key store must be in the form: safkeyring://user/keyring, where:

- safkeyring is a literal and this name is case insensitive
- user is the RACF user id that owns the keyring
- keyring is the name of the RACF keyring and the name of the keyring is case sensitive

The following example uses the standard AMS keyring for user JOHNDOE:

```
jceracfs.keystore=safkeyring://JOHNDOE/drq.ams.keyring
```

## **IBM MQ classes for Java environment-dependent behavior**

IBM MQ classes for Java allow you to create applications that can run against different versions of IBM MQ. This collection of topics describes the behavior of the Java classes dependent on these different versions.

IBM MQ classes for Java provides a core of classes, which provide consistent function and behavior in all the environments. Features outside this core depend on the capability of the queue manager to which the application is connected.

Except where noted here, the behavior exhibited is as described in the MQI application reference appropriate to the queue manager.

### **Core classes in IBM MQ classes for Java:**

IBM MQ classes for Java contains a core set of classes, which can be used in all environments.

The following set of classes are considered core classes, and can be used in all environments with only the minor variations listed in “Restrictions and variations for core classes of IBM MQ classes for Java” on page 1088.

- MQEnvironment
- MQException
- MQGetMessageOptions
  - Excluding:
    - MatchOptions
    - GroupStatus
    - SegmentStatus
    - Segmentation
- MQManagedObject
  - Excluding:
    - inquire()
    - set()
- MQMessage
  - Excluding:
    - groupId
    - messageFlags
    - messageSequenceNumber
    - offset
    - originalLength
- MQPoolServices
- MQPoolServicesEvent
- MQPoolServicesEventListener

- MQPoolToken
- MQPutMessageOptions
  - Excluding:
    - knownDestCount
    - unknownDestCount
    - invalidDestCount
    - recordFields
- MQProcess
- MQQueue
- MQQueueManager
  - Excluding:
    - begin()
    - accessDistributionList()
- MQSimpleConnectionManager
- MQTopic
- MQC

**Note:**

1. Some constants are not included in the core (see “Restrictions and variations for core classes of IBM MQ classes for Java” for details); do not use them in completely portable programs.
2. Some platforms do not support all connection modes. On these platforms, you can use only the core classes and options that relate to the supported modes.

**Restrictions and variations for core classes of IBM MQ classes for Java:**

The core classes generally behave consistently across all environments, even if the equivalent MQI calls normally have environment differences. The behavior is as if a Windows, UNIX or Linux IBM MQ queue manager is used, except for the following minor restrictions and variations.

*Restrictions for MQGMO\_\* values in IBM MQ classes for Java:*

Certain MQGMO\_\* values are not supported by all queue managers.

Use of the following MQGMO\_\* values might result in an MQException being thrown from an MQQueue.get():

```

MQGMO_SYNCPOINT_IF_PERSISTENT
MQGMO_MARK_SKIP_BACKOUT
MQGMO_BROWSE_MSG_UNDER_CURSOR
MQGMO_LOCK
MQGMO_UNLOCK
MQGMO_LOGICAL_ORDER
MQGMO_COMPLETE_MESSAGE
MQGMO_ALL_MSGS_AVAILABLE
MQGMO_ALL_SEGMENTS_AVAILABLE
MQGMO_UNMARKED_BROWSE_MSG
MQGMO_MARK_BROWSE_HANDLE
MQGMO_MARK_BROWSE_CO_OP
MQGMO_UNMARK_BROWSE_HANDLE
MQGMO_UNMARK_BROWSE_CO_OP

```

Additionally, MQGMO\_SET\_SIGNAL is not supported when used from Java.



*Restrictions for MQPMRF\_\* values in IBM MQ classes for Java:*

These are used only when putting messages to a distribution list, and are supported only by queue managers supporting distribution lists. For example, z/OS queue managers do not support distribution lists.

*Restrictions for MQPMO\_\* values in IBM MQ classes for Java:*

Certain MQPMO\_\* values are not supported by all queue managers

Use of the following MQPMO\_\* values might result in an MQException being thrown from an MQQueue.put() or an MQQueueManager.put():

MQPMO\_LOGICAL\_ORDER  
MQPMO\_NEW\_CORREL\_ID  
MQPMO\_NEW\_MESSAGE\_ID  
MQPMO\_RESOLVE\_LOCAL\_Q

*Restrictions and variations for MQCNO\_\* values in IBM MQ classes for Java:*

Certain MQCNO\_\* values are not supported.

- Automatic client reconnect is not supported by the IBM MQ classes for Java. Whatever value MQCNO\_RECONNECT\_\* you set, the connection continues to behave as if you set MQCNO\_RECONNECT\_DISABLED.
- MQCNO\_FASTPATH is ignored on queue managers that do not support MQCNO\_FASTPATH. It is also ignored by client connections.

*Restrictions for MQRO\_\* values in IBM MQ classes for Java:*

The following report options can be set.

MQRO\_EXCEPTION\_WITH\_FULL\_DATA  
MQRO\_EXPIRATION\_WITH\_FULL\_DATA  
MQRO\_COA\_WITH\_FULL\_DATA  
MQRO\_COD\_WITH\_FULL\_DATA  
MQRO\_DISCARD\_MSG  
MQRO\_PASS\_DISCARD\_AND\_EXPIRY

For more information see Report.

*Miscellaneous differences between IBM MQ classes for Java on z/OS and other platforms:*

IBM MQ for z/OS behaves differently from IBM MQ on other platforms in some areas.

#### **BackoutCount**

A z/OS queue manager returns a maximum BackoutCount of 255, even if the message has been backed out more than 255 times.

#### **Default dynamic queue prefix**

When connected to a z/OS queue manager using a bindings connection, the default dynamic queue prefix is CSQ.\*. Otherwise, the default dynamic queue prefix is AMQ.\*.

#### **MQQueueManager constructor**

Client connect is not supported on z/OS. Attempting to connect with client options results in an MQException with MQCC\_FAILED and MQRC\_ENVIRONMENT\_ERROR.

The MQQueueManager constructor might also fail with MQRC\_CHAR\_CONVERSION\_ERROR (if it fails to initialize conversion between the IBM-1047 and ISO8859-1 code pages), or MQRC\_UCS2\_CONVERSION\_ERROR (if it fails to initialize conversion between the queue manager's code page and Unicode). If your application fails with one of these reason codes,

ensure that the National Language Resources component of Language Environment is installed, and ensure that the correct conversion tables are available.

Conversion tables for Unicode are installed as part of the z/OS C/C++ optional feature. See the *z/OS C/C++ Programming Guide, SC09-4765*, for more information about enabling UCS-2 conversions.

### **Features outside the core classes of IBM MQ classes for Java:**

IBM MQ classes for Java contain certain functions that are specifically designed to use API extensions that are not supported by all queue managers. This collection of topics describes how they behave when using a queue manager that does not support them.

#### *Variations in the MQQueueManager constructor option:*

Some of the MQQueueManager constructors include an optional integer argument. Some values of this argument are not accepted on all platforms.

Where an MQQueueManager constructor include an optional integer argument, it maps onto the MQCNO options field of the MQI, and is used to switch between normal and fast path connection. This extended form of the constructor is accepted in all environments, if the only options used are MQCNO\_STANDARD\_BINDING or MQCNO\_FASTPATH\_BINDING. Any other options cause the constructor to fail with MQRC\_OPTIONS\_ERROR. The fast path option CMQC.MQCNO\_FASTPATH\_BINDING is honored only with a bindings connection to a queue manager that supports it. In other environments, it is ignored.

#### *Restrictions on the MQQueueManager.begin() method:*

This method can be used only against an IBM MQ queue manager on UNIX, Linux, or Windows systems in bindings mode. Otherwise, it fails with MQRC\_ENVIRONMENT\_ERROR.

See “JTA/JDBC coordination using IBM MQ classes for Java” on page 1073 for more details.

#### *Variations in the MQGetMessageOptions fields:*

Some queue managers do not support the Version 2 MQGMO structure, so you must set some fields to their default values.

When using a queue manager that does not support the Version 2 MQGMO structure, leave the following fields set to their default values:

- GroupStatus
- SegmentStatus
- Segmentation

Also, the MatchOptions field supports only MQMO\_MATCH\_MSG\_ID and MQMO\_MATCH\_CORREL\_ID. If you put unsupported values into these fields, the subsequent MQDestination.get() fails with MQRC\_GMO\_ERROR. If the queue manager does not support the Version 2 MQGMO structure, these fields are not updated after a successful MQDestination.get().

*Restrictions in distribution lists in IBM MQ classes for Java:*

Not all queue managers allow you to open an MQDistributionList.

The following classes are used to create distribution lists:

- MQDistributionList
- MQDistributionListItem
- MQMessageTracker

You can create and populate MQDistributionLists and MQDistributionListItems in any environment, but not all queue managers allow you to open an MQDistributionList. In particular, z/OS queue managers do not support distribution lists. Attempting to open an MQDistributionList when using such a queue manager results in MQRC\_OD\_ERROR.

*Variations in MQPutMessageOptions fields:*

If a queue manager does not support distribution lists, certain MQPMO fields are treated differently.

Four fields in the MQPMO are rendered as the following member variables in the MQPutMessageOptions class:

- knownDestCount
- unknownDestCount
- invalidDestCount
- recordFields

These fields are primarily intended for use with distribution lists. However, a queue manager that supports distribution lists also fills in the DestCount fields after an MQPUT to a single queue. For example, if the queue resolves to a local queue, knownDestCount is set to 1 and the other two count fields are set to 0.

If the queue manager does not support distribution lists, these values are simulated as follows:

- If the put() succeeds, unknownDestCount is set to 1, and the others are set to 0.
- If the put() fails, invalidDestCount is set to 1, and the others are set to 0.

The recordFields variable is used with distribution lists. A value can be written into recordFields at any time, regardless of the environment. It is ignored if the MQPutMessageOptions object is used on a subsequent MQDestination.put() or MQQueueManager.put(), rather than MQDistributionList.put().

*Restrictions in MQMD fields with IBM MQ classes for Java:*

Certain MQMD fields concerned with message segmentation should be left at their default value when using a queue manager that does not support segmentation.

The following MQMD fields are largely concerned with message segmentation:

- GroupId
- MsgSeqNumber
- Offset
- MsgFlags
- OriginalLength

If an application sets any of these MQMD fields to values other than their defaults, and then does a put() or get() on a queue manager that does not support these, the put() or get() raises an MQException with MQRC\_MD\_ERROR. A successful put() or get() with such a queue manager always leaves the MQMD fields set to their default values. Do not send a grouped or segmented message to a Java application that runs against a queue manager that does not support message grouping and segmentation.

If a Java application attempts to get() a message from a queue manager that does not support these fields, and the physical message to be retrieved is part of a group of segmented messages (that is, it has non-default values for the MQMD fields), it is retrieved without error. However, the MQMD fields in the MQMessage are not updated, the MQMessage format property is set to MQFMT\_MD\_EXTENSION, and the true message data is prefixed with an MQMDE structure that contains the values for the new fields.

### **Restrictions for IBM MQ classes for Java under CICS Transaction Server:**

In the CICS Transaction Server for z/OS environment, only the main (first) thread is allowed to issue CICS or IBM MQ calls.

Note, that IBM MQ JMS classes are not supported for use in a CICS Java application.

It is therefore not possible to share MQQueueManager or MQQueue objects between threads in this environment, or to create a new MQQueueManager on a child thread.

**z/OS** “Miscellaneous differences between IBM MQ classes for Java on z/OS and other platforms” on page 1089 identifies some restrictions and variations that apply to the IBM MQ classes for Java when running against a z/OS queue manager. Additionally, when running under CICS, the transaction control methods on MQQueueManager are not supported. Instead of issuing MQQueueManager.commit() or MQQueueManager.backout(), applications use the JCICS task synchronization methods, Task.commit() and Task.rollback(). The Task class is supplied by JCICS in the com.ibm.cics.server package.

## **Using the IBM MQ resource adapter**

The resource adapter allows applications that are running in an application server to access IBM MQ resources. It supports inbound and outbound communication.

### **What the resource adapter contains**

The Java Platform, Enterprise Edition Connector Architecture (JCA) provides a standard way of connecting applications that are running in a Java EE environment to an Enterprise Information System (EIS) such as IBM MQ or Db2. The IBM MQ resource adapter implements the JCA 1.7 interfaces and contains the IBM MQ classes for JMS. It allows JMS applications and message driven beans (MDBs), running in an application server, to access the resources of an IBM MQ queue manager. The resource adapter supports both the point-to-point domain and the publish/subscribe domain.

The IBM MQ resource adapter supports two types of communication between an application and a queue manager:

#### **Outbound communication**

An application starts a connection to a queue manager, and then sends JMS messages to JMS destinations and receives JMS messages from JMS destinations in a synchronous manner.

#### **Inbound communication**

A JMS message that arrives at a JMS destination is delivered to an MDB, which processes the message asynchronously.

The resource adapter also contains the IBM MQ classes for Java. The classes are automatically available to applications that are running in an application server that the resource adapter has been deployed into, and allow applications that are running in that application server to use the IBM MQ classes for Java API when they are accessing resources of an IBM MQ queue manager.

The use of the IBM MQ classes for Java within a Java EE environment is supported with restrictions. For information about these restrictions, see “Running IBM MQ classes for Java applications within Java EE” on page 1023.

## Which version of the resource adapter to use

The Java Platform, Enterprise Edition (Java EE) version of the application server that you are using determines the version of the resource adapter that you must use:

### Java EE 7

The IBM MQ Version 8.0 and IBM MQ Version 9.0 resource adapter supports JCA v1.7 and provides JMS 2.0 support. This resource adapter needs to be deployed within a Java EE 7 and later application server (see “IBM MQ resource adapter statement of support” on page 1094).


You can install the Version 8.0 or later resource adapter on any application server that is certified as compliant with the Java Platform, Enterprise Edition 7 specification. Using the Version 8.0 or later resource adapter, an application can connect to a IBM WebSphere MQ Version 7.0 or later queue manager using either the BINDINGS or CLIENT transport, or to a IBM WebSphere MQ Version 6.0 queue manager using the CLIENT transport only.

**Important:** The IBM MQ Version 8.0 or later resource adapter can be deployed only into an application server that supports JMS 2.0.

### Java EE 5 and Java EE 6

The IBM WebSphere MQ Version 7.5 resource adapter supports Java EE Connector Architecture (JCA) v1.5 and provides JMS 1.1 support. To provide full integration with WebSphere Application Server Liberty, the IBM WebSphere MQ Version 7.5 resource adapter is updated to APAR IC92914 from Version 7.5.0, Fix Pack 2. This resource adapter retains full compatibility with other Java EE 5 and later application servers (see WebSphere MQ resource adapter v7.1 and later statement of support).

## Using the resource adapter with WebSphere Application Server traditional

 The IBM MQ Version 9.0 resource adapter is pre-installed within WebSphere Application Server traditional Version 9.0. Therefore, there is no requirement to install a new resource adapter.

**Note:** An IBM MQ Version 9.0 resource adapter can connect in CLIENT or BINDINGS transport mode to any in-service IBM MQ queue manager.

## Using the resource adapter with WebSphere Application Server Liberty

To connect to IBM MQ from WebSphere Application Server Liberty, you must use the IBM MQ resource adapter. Since Liberty does not contain the IBM MQ resource adapter, you must obtain it separately from Fix Central. The version of the resource adapter that you use depends on the Java EE version of the application server.

For more information about how to download and install the resource adapter, see “Installing the resource adapter in Liberty” on page 1100.

**Related concepts:**

“Configuring the resource adapter for inbound communication” on page 1106

To configure inbound communication, define the properties of one or more ActivationSpec objects.

“Configuring the resource adapter for outbound communication” on page 1121

To configure outbound communication, define the properties of a ConnectionFactory object and an administered destination object.

“Using IBM MQ classes for JMS” on page 824

IBM MQ classes for Java Message Service (IBM MQ classes for JMS) is the JMS provider that is supplied with IBM MQ. As well as implementing the interfaces defined in the javax.jms package, IBM MQ classes for JMS provides two sets of extensions to the JMS API.

“Using IBM MQ classes for Java” on page 1021

Use IBM MQ in a Java environment. IBM MQ classes for Java allow a Java application to connect to IBM MQ as an IBM MQ client, or connect directly to an IBM MQ queue manager.

**Related information:**

Configuring the application server to use the latest resource adapter maintenance level

Problem determination for the IBM MQ resource adapter

**Related information for WebSphere Application Server Version 8.5.5**

Maintaining the IBM MQ resource adapter

Deploying JMS applications to the Liberty profile to use the IBM MQ messaging provider

**IBM MQ resource adapter statement of support**

The resource adapter that comes with IBM MQ Version 9.0 implements the JMS 2.0 specification. It can be deployed only into an application server that is Java Platform, Enterprise Edition 7 (Java EE 7) compliant and therefore supports JMS 2.0.

A list of certified application servers is maintained on Oracle's web site.

**Deployment within WebSphere Application Server Liberty**

WebSphere Application Server Liberty Version 8.5.5, Fix Pack 6 and later, and WebSphere Application Server Liberty Version 9.0 and later are Java EE 7 certified application servers so the IBM MQ Version 9.0 resource adapter can be deployed into them.

WebSphere Application Server Liberty has available the wmqJmsClient-1.1 feature to allow working with JMS 1.1 resource adapters and the wmqJmsClient-2.0 feature to allow working with JMS 2.0 resource adapters. The IBM MQ Version 9.0 resource adapter must be deployed with the wmqJmsClient-2.0 feature.

Information on this configuration is in the Scenario Connecting WebSphere Application Server Liberty profile to IBM MQ.

**Deployment within WebSphere Application Server traditional**

WebSphere Application Server traditional Version 9.0 is supplied with an IBM MQ Version 9.0 resource adapter already installed. The resource adapter can connect to any queue managers that are running on a supported version of IBM MQ or IBM WebSphere MQ. For more information, see “Connectivity to IBM MQ Version 8.0 and Version 9.0 queue managers” on page 1095.

The IBM MQ Version 9.0 resource adapter cannot be deployed into earlier versions of WebSphere Application Server, as these versions are not Java EE 7 certified.

## Using the resource adapter with other application servers

For all other Java EE 7 compliant application servers, problems that occur following the successful completion of the IBM MQ resource adapter Installation Verification Test (IVT) can be reported to IBM for the investigation of IBM MQ product trace and other IBM MQ diagnostic information. If the IBM MQ resource adapter IVT cannot be run successfully, any problems that are encountered are likely to be caused by incorrect deployment or incorrect resource definitions that are application server specific and the problems must be investigated by using the application server documentation and/or the support organization for that application server.

### Java Runtime

The Java Runtime (JRE) that is used to run the application server must be one that is supported with the IBM MQ Version 9.0 Client. These JREs are listed in the System Requirements for IBM MQ Version 9.0. (Select which operating system or component report you want to see then follow the **Java** link that is listed under the **Supported Software** tab.)

### Connectivity to IBM MQ Version 8.0 and Version 9.0 queue managers

The full range of JMS 2.0 functionality is available when connecting to a Version 8.0 or Version 9.0 queue manager by using the Version 9.0 resource adapter that has been deployed into a Java EE 7 certified application server. To make use of this functionality, the resource adapter needs to connect to the queue manager by using IBM MQ messaging provider normal mode. For more information, see [../com.ibm.mq.con.doc/q123360\\_.dita#q123360\\_.](#)

### Connectivity to IBM WebSphere MQ Version 7.5 or earlier queue managers

It is supported to deploy the IBM MQ Version 9.0 resource adapter into a Java EE 7 certified application server that supports JMS 2.0 and connect that resource adapter to a queue manager that is running IBM WebSphere MQ Version 7.5 or earlier. The functionality that is available is limited by the capabilities of the queue manager. For more information, see [../com.ibm.mq.con.doc/q123360\\_.dita#q123360\\_.](#)

### MQ Extensions

The JMS 2.0 specification introduces changes to how certain behaviors work. Because IBM MQ Version 8.0 or Version 9.0 implement this specification, there are changes in behavior between these two releases and previous versions of IBM WebSphere MQ. The Version 8.0 and Version 9.0 IBM MQ classes for JMS include support for the Java system property `com.ibm.mq.jms.SupportMQExtensions` that, when set to `TRUE`, causes these two releases to revert these behaviors to those of IBM WebSphere MQ Version 7.5 or earlier. The default value of the property is `FALSE`.

The IBM MQ Version 9.0 resource adapter also includes a resource adapter property called `supportMQExtensions` that has the same effect and default value as the `com.ibm.mq.jms.SupportMQExtensions` Java system property. This resource adapter property is set to `false` in the `ra.xml` by default.

If both the resource adapter property and Java system property are set, then the system property has precedence.

Note that within the resource adapter that is already deployed within WebSphere Application Server traditional Version 9.0, this property is automatically set to `TRUE` to aid migration.

For more information, see “`SupportMQExtensions` property” on page 1001.

## General issues

### Session interleaving is not supported

Some application servers provide a capability called session interleaving, where the same JMS session can be used in multiple transactions, although it is only enlisted in one at a time. The IBM MQ resource adapter does not support this capability, which can lead to the following issues:

An attempt to put a message to a MQ queue fails with reason code 2072 (MQRC\_SYNCPOINT\_NOT\_AVAILABLE).

Calls to `xa_close()` fail with reason code -3 (XAER\_PROTO), and an FDC with probe ID AT040010 is generated on the IBM MQ queue manager being accessed from the application server. For information on how to disable this capability, see your application server documentation.

### Java Transaction API (JTA) specification of how XA resources are recovered for XA transaction recovery

Section 3.4.8 of the JTA specification does not define a specific mechanism by which XA resources are re-created to perform XA transactional recovery. As such, it is up to each individual transaction manager (and, therefore, the application server) how XA resources involved in an XA transaction are recovered. It is possible that, for some application servers, the IBM MQ Version 9.0 resource adapter does not implement the application server specific mechanisms that are used to perform XA transactional recovery.

### Matching connections in a ManagedConnectionFactory

An application server can invoke the `matchManagedConnections` method on a `ManagedConnectionFactory` instance provided by the IBM MQ resource adapter. A `ManagedConnection` is returned only if the method finds one that matches both the `javax.security.auth.Subject` and `javax.resource.spi.ConnectionRequestInfo` arguments that were passed to the method by the application server.

## Limitations of the IBM MQ resource adapter

The IBM MQ resource adapter is supported on all IBM MQ platforms. However, when you use the IBM MQ resource adapter, some features of IBM MQ are unavailable or limited.

The IBM MQ resource adapter has the following limitations:

- Since IBM MQ Version 8.0, the resource adapter is a Java Platform, Enterprise Edition 7 (Java EE 7) resource adapter providing JMS 2.0 function. Consequently, the Version 8.0 or later resource adapter must be installed in a Java EE 7 or later certified application server. It can connect in client or bindings transport mode to any in-service queue manager.
- When running inside the WebSphere Application Server Liberty application server, the stabilized IBM MQ classes for Java are not supported. Within other application servers the IBM MQ classes for Java are not recommended for use. See the IBM technote *Using WebSphere MQ Java Interfaces in J2EE/JEE Environments* for details of IBM MQ classes for Java considerations within Java EE.
- When running inside the WebSphere Application Server Liberty application server on z/OS, the `wmqJmsClient-2.0` feature must be used. Generic JCA support is not possible for z/OS.
- The IBM MQ resource adapter does not support channel exit programs that are written in languages other than Java.
- While an application server is running, the value of the `sslFipsRequired` property must be true for all JCA resources or false for all JCA resources. This is a requirement even if the JCA resources are not used concurrently. If the `sslFipsRequired` property has different values for different JCA resources, IBM MQ issues the reason code `MQRC_UNSUPPORTED_CIPHER_SUITE`, even if a TLS connection is not being used.
- You cannot specify more than one keystore for an application server. If connections are made to more than one queue manager, all the connections must use the same keystore. This limitation does not apply to WebSphere Application Server.
- If you use a client channel definition table (CCDT) with more than one suitable client connection channel definition, in the event of a failure the resource adapter might select a different channel



definition and therefore a different queue manager from the CCDT, which would cause problems for transaction recovery. The resource adapter does not take any action to prevent such a configuration from being used, and it is your responsibility to avoid configurations that might cause problems for transaction recovery.

- The connection retry functionality introduced in IBM WebSphere MQ Version 7.0.1 is not supported for outbound connections when running in a Java EE container (EJB/Servlet). Connection retry is not supported at all for outbound JMS when the adapter is used in a JEE container context, regardless of transaction configuration or for non-transacted use. To configure equivalent reconnect function, see the scenario Making the messaging infrastructure highly available.
- Re-authentication, as defined in Section 9.1.9 of the Java EE Connector Architecture version 1.7 specification, of JMS connections is not supported. The `ra.xml` file within the IBM MQ resource adapter must have the property called **reauthentication-support** set to the value `false`. An attempt by the application server to re-authenticate a JMS connection results in the IBM MQ resource adapter throwing a `javax.resource.spi.SecurityException` with the `MQJCA1028` message code.

**Related information:**

Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client  
Federal Information Processing Standards (FIPS) for UNIX, Linux and Windows

### **WebSphere Application Server and the IBM MQ resource adapter**

The IBM MQ resource adapter is used by applications that perform JMS messaging with the IBM MQ messaging provider in WebSphere Application Server.

**Important:** Do not use the IBM MQ or IBM WebSphere MQ resource adapter with WebSphere Application Server Version 6.0 or Version 6.1.

WebSphere Application Server Version 7.0 and WebSphere Application Server Version 8.0 include a version of the IBM WebSphere MQ Version 7.0 resource adapter.

WebSphere Application Server Version 8.5.5 includes a version of the IBM WebSphere MQ Version 7.1 resource adapter.

**V 9.0.0** WebSphere Application Server traditional Version 9.0 includes a version of the IBM MQ Version 9.0 resource adapter. The IBM MQ Version 9.0 resource adapter cannot be deployed into earlier versions of WebSphere Application Server, as these versions are not Java EE 7 certified.

If you want to use a JMS application to access the resources of an IBM MQ queue manager from within WebSphere Application Server, use the IBM MQ messaging provider in WebSphere Application Server. The IBM MQ messaging provider contains a version of the IBM MQ classes for JMS. For more information, see Identifying which IBM MQ client is shipped with WebSphere Application Server


**Important:** Do not include any of the IBM MQ classes for JMS or IBM MQ classes for Java JAR files within your application. Doing so can result in `ClassCastException`s and can be difficult to maintain.

## Liberty and the IBM MQ resource adapter

The IBM MQ resource adapter can be installed into WebSphere Application Server Liberty Version 8.5.5, Fix Pack 2 or later, by using either the `wmqJmsClient-1.1` or `wmqJmsClient-2.0` feature, depending on which version of the resource adapter you are installing. Alternatively you can, subject to some restrictions, install the resource adapter by using generic Java Platform, Enterprise Edition Connector Architecture (Java EE JCA) support.

### General restrictions when installing the resource adapter into Liberty

The following restrictions apply to the resource adapter when using either the `wmqJmsClient-1.1` or `wmqJmsClient-2.0` feature and also when using generic JCA support:

- The IBM MQ classes for Java are not supported in Liberty. They must not be used with either the IBM MQ Liberty messaging feature or with the generic JCA support. For more information, see *Using WebSphere MQ Java Interfaces in J2EE/JEE Environments*.
- The IBM MQ resource adapter has a transport type of `BINDINGS_THEN_CLIENT`. This transport type is not supported within the IBM MQ Liberty messaging feature.
-  Before IBM MQ Version 9.0, the Advanced Message Security (AMS) feature was not included in the IBM MQ Liberty messaging feature. However, AMS is supported with an IBM MQ Version 9.0 resource adapter.

### Restrictions when using the Liberty features

With Liberty Version 8.5.5, Fix Pack 2 to Version 8.5.5, Fix Pack 5 inclusive, only the `wmqJmsClient-1.1` feature was available and only JMS 1.1 could be used. Liberty Version 8.5.5, Fix Pack 6 added the `wmqJmsClient-2.0` feature so JMS 2.0 could be used.

However, the feature that you must use depends on which version of the resource adapter you are using:

- The IBM WebSphere MQ Version 7.5.0, Fix Pack 6 and later Version 7.5 resource adapter can be used with the `wmqJmsClient-1.1` feature only.
- The IBM MQ Version 8.0.0, Fix Pack 3 and later Version 8.0 resource adapter can be used with `wmqJmsClient-2.0` feature only.
- The IBM MQ Version 9.0 resource adapter can be used with `wmqJmsClient-2.0` feature only.

### Restrictions when using generic JCA support

If you are using generic JCA support, the following restrictions apply:

- You must specify the level of JMS when using the generic JCA support:
  - JMS 1.1 and JCA 1.6 must be used only with the IBM WebSphere MQ Version 7.5.0, Fix Pack 6 and later Version 7.5 resource adapters.
  - JMS 2.0 and JCA 1.7 must be used only with the IBM MQ Version 8.0.0, Fix Pack 3 and later Version 8.0 resource adapters.
- It is not possible to run the IBM MQ resource adapter on z/OS using generic JCA support. In order to run the IBM MQ resource adapter on z/OS, it must be run with the `wmqJmsClient-1.1` or `wmqJmsClient-2.0` feature.
- Tracing, and logging are not integrated within the Liberty trace system. Trace is written to a separate file, and it must be enabled by setting the system property. This procedure is the same as configuring the IBM MQ classes for JMS trace facility for a Java Standard Environment. For more information, see “Java Standard Environment Trace stanza” on page 839.
- The location of the resource adapter is specified by using the following xml element:

```
<resourceAdapter id="mqJms" location="{server.config.dir}/wmq.jmsra.rar">
 <classloader apiTypeVisibility="spec, ibm-api, api, third-party"/>
</resourceAdapter>
```

**Important:** The value of the ID tag can be anything EXCEPT for `wmqJms`. If you do use `wmqJms` as the ID, then Liberty is not able to properly load the resource adapter. This is because `wmqJms` is the ID that is used internally to refer to the specific feature for IBM MQ. It actually creates a `NullPointerException`. The following examples show some snippets from a `server.xml` file:

```
<!-- Enable features -->
<featureManager>
 <feature>servlet-3.1</feature>
 <feature>jndi-1.0</feature>
 <feature>jca-1.7</feature>
 <feature>jms-2.0</feature>
</featureManager>
```

**Tip:** Note the use of the `jca-1.7` and `jms-2.0` features and the lack of the `wmqJmsClient-2.0` feature.

```
<resourceAdapter id="mqJms" location="{server.config.dir}/wmq.jmsra.rar">
 <classloader apiTypeVisibility="spec, ibm-api, api, third-party"/>
</resourceAdapter>
```

**Tip:** Note the use of `mqJms` for the ID, which is preferred. Do not use `wmqJms`.

```
<application id="WMQHTTP" location="{server.config.dir}/apps/WMQHTTP.war"
name="WMQHTTP" type="war">
 <classloader apiTypeVisibility="spec, ibm-api, api, third-party"
classProviderRef="mqJms"/>
</application>
```

**Tip:** Note the `classloaderProviderRef` back to the resource adapter through the id `mqJms`; this is to permit IBM MQ specific classes to be loaded.

## Installing the IBM MQ resource adapter

The IBM MQ resource adapter is supplied as a resource archive (RAR) file. Install the RAR file in your application server. You might need to add directories to the system path.

### About this task

The IBM MQ resource adapter is supplied as a resource archive (RAR) file called `wmq.jmsra.rar`. The RAR file contains IBM MQ classes for JMS and the IBM MQ implementation of the Java EE Connector Architecture (JCA) interfaces.

When you install the resource adapter as part of the IBM MQ product installation, `wmq.jmsra.rar` is installed with IBM MQ classes for JMS in the directory shown in Table 156.

Table 156. The directory containing `wmq.jmsra.rar` for each platform

Platform	Directory
UNIX and Linux	<code>MQ_INSTALLATION_PATH/java/lib/jca</code>
IBM i	<code>/QIBM/ProdData/mqm/java/lib/jca</code>
Windows	<code>MQ_INSTALLATION_PATH\java\lib\jca</code>




`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM MQ is installed.

You must use the IBM MQ resource adapter to connect to IBM MQ from an application server. Depending on which application server you are using, the resource adapter might be pre-installed, or you might need to install it yourself.

Table 157. Installation of the resource adapter in an application server

Application server	Pre-installed or need to install?
WebSphere Application Server traditional Version 9.0	The IBM MQ Version 9.0 resource adapter is pre-installed within WebSphere Application Server traditional Version 9.0. Therefore, you do not need to install a new resource adapter in WebSphere Application Server traditional Version 9.0.
WebSphere Application Server Liberty	WebSphere Application Server Liberty does not contain the IBM MQ resource adapter, so you must obtain it separately from Fix Central.
Other Java EE application server	Obtain the resource adapter separately from Fix Central, as for WebSphere Application Server Liberty.

## Procedure

- If you are connecting to IBM MQ from WebSphere Application Server Liberty, or another Java EE application server, download and install the IBM MQ resource adapter as described in “Installing the resource adapter in Liberty.”
-   For bindings connections on UNIX and Linux systems, ensure that the directory containing the Java Native Interface (JNI) libraries is in the system path. For the location of this directory, which also contains the IBM MQ classes for JMS libraries, see “Configuring the Java Native Interface (JNI) libraries” on page 835.
-  On Windows, this directory is automatically added to the system path during the installation of IBM MQ classes for JMS.

**Tip:** As an alternative to setting the system path, the IBM MQ resource adapter has a property called `nativeLibraryPath` that can be used to specify the location of the JNI library. For example in WebSphere Application Server Liberty this would be configured as shown in the following example:

```
<wmqJmsClientnativeLibraryPath="/opt/mqm/java/lib64"/>
```

Transactions are supported in both client and bindings mode.

### Installing the resource adapter in Liberty:

To connect to IBM MQ from WebSphere Application Server Liberty, or other Java EE application servers, you must use the IBM MQ resource adapter. Since Liberty does not contain the IBM MQ resource adapter, you must obtain it separately from Fix Central.

### Before you begin

**Note:** The information in this topic does not apply to WebSphere Application Server traditional Version 9.0. The IBM MQ Version 9.0 resource adapter is preinstalled within WebSphere Application Server traditional Version 9.0. Therefore, there is no requirement to install a new resource adapter in this case.

Before you start this task, make sure that you have a Java runtime environment (JRE) installed on your machine and that the JRE has been added to the system path.

The Java installer that is used in this installation process does not require running as root or any specific user. The only requirement is that the user it is run as has access write to the directory that you want the files to go in.

## About this task

The JAR file for the resource adapter that you can download from Fix Central is executable. When you run this executable file, it displays the IBM MQ license agreement, which must be accepted. It asks for a directory in which to install the IBM MQ resource adapter. The resource adapter RAR file and installation verification test (IVT) program are then installed in that directory. You can either accept the default or specify another directory, which might be the resource adapters directory of an application server, or any other directory on your system. The directory is created as part of the installation if it does not exist.

Before IBM MQ Version 9.0, the name of the file to be downloaded was in the format of *V.R.M.F-WS-MQ-Java-InstallRA.jar*, for example *8.0.0.6-WS-MQ-Java-InstallRA.jar*. From IBM MQ Version 9.0, the format of the file name is *V.R.M.F-IBM-MQ-Java-InstallRA.jar*, for example *9.0.0.0-IBM-MQ-Java-InstallRA.jar*.

After you have downloaded and installed the resource adapter, you are ready to configure it in WebSphere Application Server Liberty.

## Procedure

1. Download the IBM MQ resource adapter from Fix Central:
  - a. Click **Find Product** then add the information for your IBM MQ installation to the following fields:
    - In the **Product Selector** field, type MQ then select **WebSphere MQ** from the displayed list.
    - In the **Installed Version** field, click the arrow then select the version number from the displayed list, for example **9.0.0.0**.
    - In the **Platform** field, click the arrow and select your platform, for example, **Windows 64-bit, x86**.

Click **Continue**.

- b. Make sure that **Browse for Fixes** is selected and, under **Additional Query options**, clear **Show fixes that apply to this version**, and select **Show fixes that get me to this version**, then click **Continue**. Fix Central searches for the available fixes for your selected product, version and platform, for example **WebSphere, WebSphere MQ (9.0.0.0, Windows 64-bit, x86)**.
- c. Find the resource adapter in the displayed list of available fixes. For example:

```
release level: 9.0.0.0-IBM-MQ-Install-Java-All
9.0.0.0 MQ Resource Adapter for use with Application Servers
```

Then click the resource adapter file name and follow the download process.

2. Start the installation by entering the following command from the directory to which you downloaded the file. From IBM MQ Version 9.0, the format of the command is as follows:  
`java -jar V.R.M.F-IBM-MQ-Java-InstallRA.jar`

where *V.R.M.F* is the Version, Release, Modification, and Fix Pack number and *V.R.M.F-IBM-MQ-Java-InstallRA.jar* is the name of the file that was downloaded from Fix Central. For example, to install the IBM MQ resource adapter for the Version 9.0.0.0 release, you would use the following command:

```
java -jar 9.0.0.0-IBM-MQ-Java-InstallRA.jar
```

**Note:** To carry out this installation, you must have a JRE installed on your machine and added to the system path.

When you enter the command, the following information is displayed:

```
Before you can use, extract, or install IBM MQ V9.0, you must accept
the terms of 1. IBM International License Agreement for Evaluation of
Programs 2. IBM International Program License Agreement and additional
license information. Please read the following license agreements carefully.
```

The license agreement is separately viewable using the --viewLicenseAgreement option.  
Press Enter to display the license terms now, or 'x' to skip.

3. Review and accept the license terms:

- a. To display the license, press Enter. Alternatively, pressing x skips the display of the license. After display of the license or immediately after selecting x, the following message appears to tell you that you can choose to display additional license terms:

Additional license information is separately viewable using the --viewLicenseInfo option.  
Press Enter to display additional license information now, or 'x' to skip.

- b. To display the additional license terms, press Enter. Alternatively, pressing x skips the display of the additional license terms. After display of the additional license terms or immediately after selecting x, the following message is displayed asking you to accept the license agreement:

By choosing the "I Agree" option below, you agree to the terms of the license agreement and non-IBM terms, if applicable. If you do not agree, select "I do not Agree".

Select [1] I Agree, or [2] I do not Agree:

- c. To accept the license agreement and continue with selecting the installation directory, select 1. Alternatively, if you select 2 the installation terminates immediately. If you selected 1, the following message appears, asking you to select a target installation directory:

Enter directory for product files or leave blank to accept the default value.  
The default target directory is H:\Liberty\WMQ  
Target directory for product files?

4. Specify the installation directory for the resource adapter:

- If you want to install the resource adapter in the default location, press Enter without specifying a value.
- If you want to install the resource adapter in a different location from the default, specify the name of the directory in which you want to install the resource adapter and then press Enter.

After the files have been installed in the selected location, a confirmation message is displayed as shown in the following example:

```
Extracting files to H:\Liberty\WMQ\wmq
Successfully extracted all product files.
```

During the installation, a new directory with the name wmq is created within the selected installation directory, and the following files are then installed in the wmq directory:

- The installation verification test program, wmq.jmsra.ivtt.
- The IBM MQ RAR file, earwmq.jmsra.rar.

5. Configure the resource adapter in WebSphere Application Server Liberty. The steps that you must take to configure the resource adapter in Liberty are as follows. For more information, see the WebSphere Application Server product documentation.

- a. Add the wmqJmsClient-2.0 feature to the server.xml file to allow working with the IBM MQ Version 9.0 resource adapter. The feature that you add (wmqJmsClient-1.1 or wmqJmsClient-2.0) depends on which version of the resource adapter you have installed. The IBM MQ Version 9.0 resource adapter must be deployed with the wmqJmsClient-2.0 feature. For more information, see "Which version of the resource adapter to use" on page 1093.
- b. Add a reference to the wmq.jmsra.rar file that you have installed.

**Note:** For Liberty versions up to Version 8.5.5, Fix Pack 1, if an EJB is deployed using solely the configuration within the ejb-jar.xml, the version of WebSphere Application Server that the Liberty Profile is using must have APAR PM89890 applied. This method of configuration is used for the resource adapter's installation verification program (IVT), so this APAR is required in order for the IVT to run.

An example configuration to support servlets and MDBs, with JNDI might look like this:

```

<featureManager>
 <feature>wmqJmsClient-1.1</feature>
 <feature>servlet-3.0</feature>
 <feature>jmsMdb-3.1</feature>
 <feature>jndi-1.0</feature>
</featureManager>

<variable name="wmqJmsClient.rar.location"
 value="H:\Liberty\WMQ\wmq\wmq.jmsra.rar"/>

```

## Configuring the IBM MQ resource adapter

To configure the IBM MQ resource adapter, you define various Java Platform, Enterprise Edition Connector Architecture (JCA) resources and, optionally, system properties. You must also configure the resource adapter to run the installation verification test (IVT) program. This is important because IBM service might require this program to be run to indicate that any non-IBM application server has been correctly configured.

### Before you begin

This task assumes that you are already familiar with JMS and IBM MQ classes for JMS. Many of the properties used to configure the IBM MQ resource adapter are equivalent to properties of IBM MQ classes for JMS objects and have the same function.

### About this task

Every application server provides its own set of administration interfaces. Some application servers provide graphical user interfaces to define JCA resources, but others require the administrator to write XML deployment plans. It is therefore beyond the scope of this documentation to provide information about how to configure the IBM MQ resource adapter for each application server.

The following steps therefore focus only on what you need to configure. Refer to documentation supplied with your application server for information about how to configure a JCA resource adapter.

### Procedure

Define JCA resources in the following categories:

- Define the properties of the ResourceAdapter object. These properties, which represent the global properties of the resource adapter, such as the level of diagnostic tracing, are described in “Configuration for ResourceAdapter object properties” on page 1104.
- Define the properties of an ActivationSpec object. These properties determine how an MDB is activated for inbound communication. For more information, see “Configuring the resource adapter for inbound communication” on page 1106.
- Define the properties of a ConnectionFactory object. The application server uses these properties to create a JMS ConnectionFactory object for outbound communication. For more information, see “Configuring the resource adapter for outbound communication” on page 1121.
- Define the properties of an administered destination object. The application server uses these properties to create a JMS Queue object or JMS Topic object for outbound communication. For more information, see “Configuring the resource adapter for outbound communication” on page 1121.
- Optional: Define a deployment plan for the resource adapter. The IBM MQ resource adapter RAR file contains a file called META-INF/ra.xml, which contains a deployment descriptor for the resource adapter. This deployment descriptor is defined by the XML schema at [http://xmlns.jcp.org/xml/ns/javaee/connector\\_1\\_7.xsd](http://xmlns.jcp.org/xml/ns/javaee/connector_1_7.xsd) and contains information about the resource adapter and the services that it provides. An application server might also require a deployment plan for the resource adapter. This deployment plan is specific to the application server.

Specify JVM system properties as required:

- If you are using Transport Layer Security (TLS), specify the locations of the keystore file and truststore file as JVM system properties, as in the following example:

```
java ... -Djavax.net.ssl.keyStore= key_store_location
 -Djavax.net.ssl.trustStore=trust_store_location
 -Djavax.net.ssl.keyStorePassword=key_store_password
```

These properties cannot be properties of an `ActivationSpec` or `ConnectionFactory` object, and you cannot specify more than one keystore for an application server. The properties apply to the whole JVM, and might therefore affect the application server if other applications, running in the application server, are using TLS connections. The application server might also reset these properties to different values. For more information about using TLS with IBM MQ classes for JMS, see “Using TLS with IBM MQ classes for JMS” on page 971.

- Optional: If required, configure the resource adapter to log warning messages to your application server's standard output log. The resource adapter logs, warning, and error messages use the same mechanism as the IBM MQ classes for JMS. For more information, see “Logging and IBM MQ classes for JMS” on page 863. This means that, by default, the messages go to a file called `mqjms.log`. To configure the resource adapter to additionally log warning messages to your application server's standard output log, set the following JVM system property for your application server:

```
-Dcom.ibm.msg.client.commonservices.log.outputName=mqjms.log,stdout
```

This is the same property as the one that is used to control trace for the IBM MQ classes for JMS. As with IBM MQ classes for JMS, it is possible to use a system property pointing to the `jms.config` file (see “The IBM MQ classes for JMS configuration file” on page 837). For information on how to set a JVM system property, see your application server documentation.

Configure the resource adapter to run the installation verification test

- Configure the resource adapter to run the installation verification test (IVT) program supplied with the IBM MQ resource adapter. For information about what you need to configure in order to run the IVT program, see “Verifying the resource adapter installation” on page 1135.

This is important because IBM service might require this program to be run to indicate that any non-IBM application server has been correctly configured.

**Important:** You must configure the resource adapter before you can run the program.

### Configuration for ResourceAdapter object properties:

The `ResourceAdapter` object encapsulates the global properties of the IBM MQ resource adapter, such as the level of diagnostic tracing. To define these properties, use the facilities of your resource adapter, as described in the documentation supplied with your application server.

The `ResourceAdapter` object has two sets of properties:

- Properties associated with diagnostic tracing
- Properties associated with the connection pool managed by the resource adapter




The way in which you define these properties depends on the administration interfaces that your application server provides. If you are using WebSphere Application Server traditional, see “WebSphere Application Server traditional configuration” on page 1106 or if you are using WebSphere Application Server Liberty, see “WebSphere Application Server Liberty configuration” on page 1106. For other application servers, see the product documentation for your application server.

For more information about defining properties associated with diagnostic trace, see [Tracing the IBM MQ Resource Adapter](#)

The resource adapter manages an internal connection pool of JMS connections that are used to deliver messages to MDBs. Table 158 on page 1105 lists the properties of the `ResourceAdapter` object that are associated with the connection pool.



Table 158. Properties of the ResourceAdapter object that are associated with the connection pool

Name of property	Type	Default value	Description
maxConnections	String	50	The maximum number of connections to an IBM MQ queue manager and the maximum number of MDBs deployed.
connectionConcurrency	String	1	The maximum number of MDBs to share a JMS connection. Sharing connections is not possible and this property always has the value 1.
reconnectionRetryCount	String	5	The maximum number of attempts made by the resource adapter to reconnect to an IBM MQ queue manager if a connection fails.
reconnectionRetryInterval	String	300 000	The time, in milliseconds, that the resource adapter waits before trying to reconnect to an IBM MQ queue manager.
startupRetryCount	String	0	The default number of times to try and connect a MDB on startup, if the queue manager is not running when the application server is started.
startupRetryInterval	String	30 000	The default sleep time between startup connection attempts (in milliseconds).
 supportMQExtensions	String	false	Reverts the IBM MQ JMS behavior to pre-JMS 2.0 behavior. For more information, see "SupportMQExtensions property" on page 1001.
 nativeLibraryPath	String	<empty>	The path to use to load the IBM MQ JNI library to permit bindings mode connections.   <b>Windows</b> On Windows the system path also needs to contain the location of the matching IBM MQ installation.

When an MDB is deployed in the application server, a new JMS connection is created and a conversation started with the queue manager, provided the maximum number of connections specified by the maxConnection property is not exceeded. The maximum number of MDBs therefore equals the maximum number of connections. If the number of deployed MDBs reaches this maximum, any attempt to deploy another MDB fails. If an MDB is stopped, its connection can be used by another MDB.

In general, if many MDBs are to be deployed, you must increase the value of the maxConnections property.

The reconnectionRetryCount and reconnectionRetryInterval properties govern the behavior of the resource adapter when connections to an IBM MQ queue manager fail, because of a network failure for example. When a connection fails, the resource adapter suspends the delivery of messages to all MDBs supplied by that connection for an interval specified by the reconnectionRetryInterval property. The resource adapter then attempts to reconnect to the queue manager. If the attempt fails, the resource adapter makes further attempts to reconnect at intervals specified by the reconnectionRetryInterval property until the limit imposed by the reconnectionRetryCount property is reached. If all attempts fail, delivery is stopped permanently until the MDBs are restarted manually.

In general, the ResourceAdapter object requires no administration. However, to enable diagnostic tracing on UNIX and Linux systems for example, you can set the following properties:

```
traceEnabled: true
traceLevel: 10
```

These properties have no effect if the resource adapter has not been started, which is the case, for example, when applications using IBM MQ resources are running only in the client container. In this situation, you can set the properties for diagnostic tracing as Java Virtual Machine (JVM) system properties. You can set the properties by using the `-D` flag on the `java` command, as in the following example:

```
java ... -DtraceEnabled=true -DtraceLevel=6
```

You do not need to define all the properties of the ResourceAdapter object. Any properties left unspecified take their default values. In a managed environment, it is better not to mix the two ways of specifying properties. If you do mix them, the JVM system properties take precedence over the properties of the ResourceAdapter object.

### WebSphere Application Server traditional configuration

The same properties are available for the resource adapter in WebSphere Application Server traditional, but they should be set for within the properties panel of the resource adapter (see JMS provider settings in the WebSphere Application Server traditional product documentation). Trace is controlled by the diagnostics section of the WebSphere Application Server traditional configuration. For more information, see Working with Diagnostic Providers in the WebSphere Application Server traditional product documentation.

### WebSphere Application Server Liberty configuration

The resource adapter is configured using XML elements in the `server.xml` file, as shown in the following example:

```
<featureManager>
...
 <feature>wmqJmsClient-2.0</feature>
...
</featureManager>
 <variable name="wmqJmsClient.rar.location"
 value="F:/_rtc_wmq8005/_build/ship/lib/jca/wmq.jmsra.rar"/>
...
 <wmqJmsClient supportMQExtensions="true" logWriterEnabled="true"/>
```

Trace is enabled by adding this XML element:

```
<logging traceSpecification="JMSApi=all:WAS.j2c=all:"/>
```

### Configuring the resource adapter for inbound communication:

To configure inbound communication, define the properties of one or more ActivationSpec objects.

The properties of an ActivationSpec object determine how a message drive bean (MDB) receives JMS messages from an IBM MQ queue. The transactional behavior of the MDB is defined in its deployment descriptor.

An ActivationSpec object has two sets of properties:

- Properties that are used to create a JMS connection to an IBM MQ queue manager
- Properties that are used to create a JMS connection consumer that delivers messages asynchronously as they arrive on a specified queue

The way in which you define the properties of an ActivationSpec object depends on the administration interfaces provided by your application server.

## New ActivationSpec properties in JMS 2.0

The JMS 2.0 specification introduced two new ActivationSpec properties. The connectionFactoryLookup and destinationLookup properties can be provided with a JNDI name of an administered object to be used in preference to the other ActivationSpec properties.

For example, assume that a connection factory is defined in JNDI and the JNDI name of that object is specified in the connectionFactoryLookup property for an activation specification. All the properties of the connection factory that are defined in JNDI are used in preference to the properties in Table 159.

If a destination is defined in JNDI and the JNDI name is set in the ActivationSpec's destinationLookup property then the values of that are used in preference to the values in Table 160 on page 1115. For more information about how these two properties are used, see "ActivationSpec connectionFactoryLookup and destinationLookup properties" on page 1117.

### Properties used to create a JMS connection to an IBM MQ queue manager

All of the properties in Table 159 are optional.

Table 159. Properties of an ActivationSpec object that are used to create a JMS connection

Name of property	Type	Valid values (default value in bold)	Description
applicationName	String	<ul style="list-style-type: none"> <li>The invoking class name, if it is available, adjusted to be no longer than 28 characters. If it is not available, the string WebSphere MQ Client for Java is used.</li> </ul>	The name by which an application is registered with the queue manager. This application name is shown by the <b>DISPLAY CONN MQSC/PCF</b> command (where the field is called <b>APPLTAG</b> ) or in the IBM MQ Explorer Application Connections display (where the field is called <b>App name</b> ).
brokerCCDurSubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li><b>SYSTEM.JMS.D.CC.SUBSCRIBER.QUEUE</b></li> <li>A queue name</li> </ul>	The name of the queue from which a connection consumer receives durable subscription messages
brokerCCSubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li><b>SYSTEM.JMS.ND.CC.SUBSCRIBER.QUEUE</b></li> <li>A queue name</li> </ul>	The name of the queue from which a connection consumer receives nondurable subscription messages
brokerControlQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li><b>SYSTEM.BROKER.CONTROL.QUEUE</b></li> <li>A queue name</li> </ul>	The name of the broker control queue
brokerQueueManager <sup>1</sup>	String	<ul style="list-style-type: none"> <li>"" (<b>empty string</b>)</li> <li>A queue manager name</li> </ul>	The name of the queue manager on which the broker is running
brokerSubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li><b>SYSTEM.JMS.ND.SUBSCRIBER.QUEUE</b></li> <li>A queue name</li> </ul>	The name of the queue from which a nondurable message consumer receives messages

Table 159. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
brokerVersion <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>unspecified</b> - After the broker is migrated from V6 to V7, set this property so that RFH2 headers are no longer used. After migration, this property is no longer relevant.</li> <li>• <b>V1</b> - To use an IBM MQ publish/subscribe broker. This value is the default value if TRANSPORT is set to BIND or CLIENT.</li> <li>• <b>V2</b> - To use a broker of IBM Integration Bus in native mode. This value is the default value if TRANSPORT is set to DIRECT or DIRECTHTTP.</li> </ul>	The version of the broker being used
ccdtURL	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A uniform resource locator (URL)</li> </ul>	A URL that identifies the name and location of the file containing the client channel definition table (CCDT) and specifies how the file can be accessed
CCSID	String	<ul style="list-style-type: none"> <li>• <b>819</b></li> <li>• A coded character set identifier supported by the Java virtual machine (JVM)</li> </ul>	The coded character set identifier for a connection
channel	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.DEF.SVRCONN</b></li> <li>• The name of an MQI channel</li> </ul>	The name of the MQI channel to use
cleanupInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>3 600 000</b></li> <li>• A positive integer</li> </ul>	The interval, in milliseconds, between background runs of the publish/subscribe cleanup utility
cleanupLevel <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>SAFE</b></li> <li>• NONE</li> <li>• STRONG</li> <li>• FORCE</li> <li>• NONDUR</li> </ul>	The cleanup level for a broker-based subscription store
clientID	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A client identifier</li> </ul>	The client identifier for a connection
cloneSupport	String	<ul style="list-style-type: none"> <li>• <b>DISABLED</b> - Only one instance of a durable topic subscriber can run at a time.</li> <li>• <b>ENABLED</b> - Two or more instances of the same durable topic subscriber can run simultaneously, but each instance must run in a separate Java virtual machine (JVM).</li> </ul>	Whether two or more instances of the same durable topic subscriber can run simultaneously

Table 159. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
connectionFactoryLookup	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• The JNDI name for a ConnectionFactory object</li> </ul>	<p>If this property is set, the ActivationSpec looks up a JMS ConnectionFactory object with the specified JNDI name in the JNDI namespace of the application server, and then uses the properties of that object to create a JMS connection to an IBM MQ queue manager, with one exception. The only property of the ActivationSpec that will be used when creating the JMS connection is the clientID. For more information, see “ActivationSpec connectionFactoryLookup and destinationLookup properties” on page 1117.</p>
connectionNameList	String	<ul style="list-style-type: none"> <li>• <b>localhost(1414)</b></li> <li>• A string composed of items separated by commas where each item takes the format: <i>HOSTNAME(PORT)</i></li> </ul> <p>where <i>HOSTNAME</i> is either a DNS name or an IP address.</p>	<p>A list of TCP/IP connection names used for inbound communications.</p> <p>When specified, <b>connectionNameList</b> supersedes the <b>hostname</b> and <b>port</b> properties.</p> <p>This property is used to reconnect to multi-instance queue managers.</p> <p><b>connectionNameList</b> is similar in form to <b>localAddress</b>, but must not be confused with it. <b>localAddress</b> specifies the characteristics of the local communications, whereas <b>connectionNameList</b> specifies how to reach a remote queue manager.</p>
failIfQuiesce	Boolean	<ul style="list-style-type: none"> <li>• <b>true</b></li> <li>• false</li> </ul>	<p>Whether calls to certain methods fail if the queue manager is in a quiescing state</p>
headerCompression	String	<ul style="list-style-type: none"> <li>• <b>NONE</b></li> <li>• SYSTEM - RLE message header compression is performed</li> </ul>	<p>A list of the techniques that can be used for compressing header data on a connection</p>

Table 159. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
hostName	String	<ul style="list-style-type: none"> <li>• <b>localhost</b></li> <li>• A host name</li> <li>• An IP address</li> </ul>	<p>The host name or IP address of the system on which the queue manager resides.</p> <p>The <b>hostname</b> and <b>port</b> properties are superseded by the <b>connectionNameList</b> property when it is specified.</p>
localAddress	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string in the format:  <code>[ host_name ][(low_port [, high_port ])]</code>                       where <i>host_name</i> is a host name or IP address, <i>low_port</i> and <i>high_port</i> are TCP port numbers, and brackets denote an optional component</li> </ul>	<p>For a connection to a queue manager, this property specifies either or both of the following things:</p> <ul style="list-style-type: none"> <li>• The local network interface to be used</li> <li>• The local port, or range of local ports, to be used</li> </ul> <p><b>LocalAddress</b> is similar in form to <b>connectionNameList</b>, but must not be confused with it. <b>LocalAddress</b> specifies the characteristics of the local communications, whereas <b>connectionNameList</b> specifies how to reach a remote queue manager.</p>
messageCompression	String	<ul style="list-style-type: none"> <li>• <b>NONE</b></li> <li>• A list of one or more of the following values separated by blank characters:                      RLE                      ZLIBFAST                      ZLIBHIGH</li> </ul>	<p>A list of the techniques that can be used for compressing message data on a connection</p>
messageRetention <sup>1</sup>	Boolean	<ul style="list-style-type: none"> <li>• <b>true</b> - Unwanted messages remain on the input queue</li> <li>• <b>false</b> - Unwanted messages are dealt with according to their disposition options</li> </ul>	<p>Whether the connection consumer keeps unwanted messages on the input queue</p>
messageSelection <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>CLIENT</b></li> <li>• <b>BROKER</b></li> </ul>	<p>Determines whether message selection is done by IBM MQ classes for JMS or by the broker. Message selection by the broker is not supported when <code>brokerVersion</code> has the value 1.</p>
password	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A password</li> </ul>	<p>The default password to use when creating a connection to the queue manager</p>

Table 159. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
pollingInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>5000</b></li> <li>• Any positive integer</li> </ul>	If each message listener within a session has no suitable message on its queue, this value is the maximum interval, in milliseconds, that elapses before each message listener tries again to get a message from its queue. If it frequently happens that no suitable message is available for any of the message listeners in a session, consider increasing the value of this property. This property is relevant only if TRANSPORT has the value BIND or CLIENT.
port	int	<ul style="list-style-type: none"> <li>• <b>1414</b></li> <li>• A TCP port number</li> </ul>	The port on which the queue manager listens.  The <b>hostname</b> and <b>port</b> properties are superseded by the <b>connectionNameList</b> property when it is specified.
providerVersion	string	<ul style="list-style-type: none"> <li>• <b>unspecified</b></li> <li>• A string in one of the following formats                             <ul style="list-style-type: none"> <li>– V.R.M.F</li> <li>– V.R.M</li> <li>– V.R</li> <li>– V</li> </ul> </li> </ul> <p>where V, R, M, and F are integer values greater than or equal to zero.</p>	The version, release, modification level and fix pack of the queue manager to which the MDB intends to connect.
queueManager	String	<ul style="list-style-type: none"> <li>• "" (<b>empty string</b>)</li> <li>• A queue manager name</li> </ul>	The name of the queue manager to connect to
receiveExit <sup>3</sup>	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string comprising one or more items separated by commas, where each item is the fully qualified name of a class that implements the IBM MQ classes for Java interface, MQReceiveExit</li> </ul>	Identifies a channel receive exit program, or a sequence of receive exit programs to be run in succession
receiveExitInIt	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string comprising one or more items of user data separated by commas</li> </ul>	The user data that is passed to channel receive exit programs when they are called

Table 159. Properties of an ActivationSpec object that are used to create a JMS connection (continued)


Name of property	Type	Valid values (default value in bold)	Description
rescanInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>5000</b></li> <li>• Any positive integer</li> </ul>	<p>When a message consumer in the point-to-point domain uses a message selector to select which messages it wants to receive, IBM MQ classes for JMS searches the IBM MQ queue for suitable messages in the sequence determined by the <b>MsgDeliverySequence</b> attribute of the queue. When IBM MQ classes for JMS finds a suitable message and delivers it to the consumer, IBM MQ classes for JMS resumes the search for the next suitable message from its current position in the queue. IBM MQ classes for JMS continues to search the queue in this way until it reaches the end of the queue, or until the interval of time in milliseconds, as determined by the value of this property, has expired. In each case, IBM MQ classes for JMS returns to the beginning of the queue to continue its search, and a new time interval commences.</p>
securityExit <sup>3</sup>	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• The fully qualified name of a class that implements the IBM MQ classes for Java interface, MQSecurityExit</li> </ul>	Identifies a channel security exit program
securityExitInit	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string of user data</li> </ul>	The user data that is passed to a channel security exit program when it is called
sendExit <sup>3</sup>	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string comprising one or more items separated by commas, where each item is the fully qualified name of a class that implements the IBM MQ classes for Java interface, MQSendExit</li> </ul>	Identifies a channel send exit program, or a sequence of send exit programs to be run in succession
sendExitInit	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string comprising one or more items of user data separated by commas</li> </ul>	The user data that is passed to channel send exit programs when they are called



Table 159. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
shareConvAllowed	Boolean	<ul style="list-style-type: none"> <li>• <b>NO</b> - A client connection cannot share its socket.</li> <li>• YES - A client connection can share its socket.</li> </ul>	Whether a client connection can share its socket with other top-level JMS connections from the same process to the same queue manager, if the channel definitions match
sparseSubscriptions <sup>1</sup>	Boolean	<ul style="list-style-type: none"> <li>• <b>false</b> - Subscriptions receive frequent matching messages.</li> <li>• true - Subscriptions receive infrequent matching messages. This value requires that the subscription queue can be opened for browse.</li> </ul>	Controls the message retrieval policy of a TopicSubscriber object
sslCertStores	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string of one or more LDAP URLs separated by blanks. Each LDAP URL has the format:  <code>ldap://host_name [: port ]</code>                       where <i>host_name</i> is a host name or IP address, <i>port</i> is a TCP port number, and brackets denote an optional component.</li> </ul>	The Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs) for use on a TLS connection
sslCipherSuite	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• The name of a CipherSuite</li> </ul>	The CipherSuite to use for a TLS connection
sslFipsRequired <sup>2</sup>	Boolean	<ul style="list-style-type: none"> <li>• <b>false</b></li> <li>• true</li> </ul>	Whether a TLS connection must use a CipherSuite that is supported by the IBM Java JSSE FIPS provider (IBMJSSEFIPS)
sslPeerName	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A template for distinguished names</li> </ul>	For a TLS connection, a template that is used to check the distinguished name in the digital certificate provided by the queue manager
sslResetCount	int	<ul style="list-style-type: none"> <li>• <b>0</b></li> <li>• An integer in the range 0 - 999 999 999</li> </ul>	The total number bytes sent and received by a TLS connection before the secret keys used by TLS are renegotiated
sslSocketFactory	String	A string representing the fully qualified class name of a class providing an implementation of the javax.net.ssl.SSLSocketFactory interface. Optionally including an argument to be passed to the constructor method, enclosed in parentheses.	Any connections established in the scope of the administered object use sockets obtained from this implementation of the SSLSocketFactory interface.

Table 159. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
statusRefreshInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>60000</b></li> <li>• Any positive integer</li> </ul>	The interval, in milliseconds, between refreshes of the long running transaction that detects when a subscriber loses its connection to the queue manager. This property is relevant only if <b>subscriptionStore</b> has the value QUEUE.
subscriptionStore <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>BROKER</b></li> <li>• MIGRATE</li> <li>• QUEUE</li> </ul>	Determines where IBM MQ classes for JMS stores persistent data about active subscriptions
transportType	String	<ul style="list-style-type: none"> <li>• <b>CLIENT</b></li> <li>• BINDINGS</li> <li>• BINDINGS_THEN_CLIENT</li> </ul>	<p>Whether a connection to a queue manager uses client mode or bindings mode. If the value BINDINGS_THEN_CLIENT is specified, the resource adapter first tries to make a connection in bindings mode. If this connection attempt fails the resource adapter then tries to make a client mode connection.</p> <p> <b>z/OS</b> If an activation specification that is running on a WebSphere Application Server for z/OS system has been configured to use the BINDINGS_THEN_CLIENT transport mode and a previously established connection is broken, then any reconnection attempts by the activation specification first attempt to use the BINDINGS transport mode. If the BINDINGS transport mode connection attempt is unsuccessful, the activation specification subsequently attempts a CLIENT transport mode connection.</p>
username	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A user name</li> </ul>	The default user name to use when creating a connection to a queue manager
wildcardFormat	String	<ul style="list-style-type: none"> <li>• CHAR- Recognizes character wildcards only, as used in broker version 1</li> <li>• <b>TOPIC</b> - Recognizes topic level wildcards only, as used in broker version 2</li> </ul>	Which version of wildcard syntax is to be used

**Notes:**

1. This property can be used with Version 7.0 of IBM MQ classes for JMS. It does not affect an application connected to a Version 7.0 queue manager unless the **providerVersion** property is set to a version number less than 7.
2. For important information about using the `sslFipsRequired` property, see “Limitations of the IBM MQ resource adapter” on page 1096.
3. For information on how to configure the resource adapter so that it can locate an exit, see “Configuring IBM MQ classes for JMS to use channel exits” on page 979.

### Properties used to create a JMS connection consumer

**Note:** The **destination** and **destinationType** must be defined explicitly. All the other properties in Table 160 are optional.

Table 160. Properties of an `ActivationSpec` object that are used to create a JMS connection consumer

Name of property	Type	Valid values (default value in bold)	Description
<code>destination</code>	String	A destination name	The destination from which to receive messages. The <b>useJNDI</b> property determines how the value of this property is interpreted.
<code>destinationLookup</code>	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• The JNDI name for a Destination object</li> </ul>	If this property is set, the <code>ActivationSpec</code> looks up a JMS Destination object with the specified JNDI name in the JNDI namespace of the application server, and then uses the properties of that object to create a JMS connection consumer, in preference to the other properties specified on the <code>ActivationSpec</code> . For more information, see “ <code>ActivationSpec</code> <code>connectionFactoryLookup</code> and <code>destinationLookup</code> properties” on page 1117.
<code>destinationType</code>	String	<ul style="list-style-type: none"> <li>• <code>javax.jms.Queue</code></li> <li>• <code>javax.jms.Topic</code></li> </ul>	The type of destination, a queue, or a topic
<code>maxMessages</code>	int	<ul style="list-style-type: none"> <li>• <b>1</b></li> <li>• A positive integer</li> </ul>	The maximum number of messages that can be assigned to a server session at one time. If the activation spec is delivering messages to an MDB in an XA transaction, a value of 1 is used regardless of the setting of this property.
<code>maxPoolDepth</code>	int	<ul style="list-style-type: none"> <li>• <b>10</b></li> <li>• A positive integer</li> </ul>	The maximum number of server sessions in the server session pool used by the connection consumer
<code>messageSelector</code>	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• An SQL92 message selector expression</li> </ul>	A message selector expression specifying which messages are to be delivered
<code>nonASFTimeout</code>	int	<ul style="list-style-type: none"> <li>• <b>0</b></li> <li>• A positive integer</li> </ul>	<p>A positive value indicates that non-ASF delivery is used. The value is the time, in milliseconds, that a get request waits for messages that might not have yet arrived (a get with wait call). The default value, 0, indicates that ASF delivery is used.</p> <p>This parameter is valid only when the application is running on WebSphere Application Server Version 7.0 or later.</p>

Table 160. Properties of an ActivationSpec object that are used to create a JMS connection consumer (continued)

Name of property	Type	Valid values (default value in bold)	Description
nonASFRollbackEnabled	Boolean	<ul style="list-style-type: none"> <li><b>false</b> - The message is consumed even if the MDB fails</li> <li>true - Failure within the MDB causes the message to roll back to the queue.</li> </ul>	Whether message delivery is within an IBM MQ syncpoint if the MDB is non-transacted. Ignored if the MDB is transacted or if <b>nonASFTimeout</b> is set to 0.
poolTimeout	int	<ul style="list-style-type: none"> <li><b>300000</b></li> <li>A positive integer</li> </ul>	The time, in milliseconds, that an unused server session is held open in the server session pool before being closed due to inactivity
readAheadAllowed	int	<ul style="list-style-type: none"> <li><b>DESTINATION</b> - Determine whether read ahead is allowed by referring to the queue or topic definition.</li> <li>DISABLED - Read ahead is not allowed.</li> <li>ENABLED - Read ahead is allowed.</li> <li>QUEUE - Determine whether read ahead is allowed by referring to the queue definition.</li> <li>TOPIC - Determine whether read ahead is allowed by referring to the topic definition.</li> </ul>	Whether the MDB is allowed to use read ahead to get nonpersistent messages from the destination into an internal buffer before receiving them
readAheadClosePolicy	int	<ul style="list-style-type: none"> <li><b>ALL</b> - All messages in the internal read ahead buffer are delivered to the MDB before it stops.</li> <li>CURRENT - Only the current MDB invocation completes, potentially leaving messages in the internal read ahead buffer, which are then discarded.</li> </ul>	What happens to messages in the internal read ahead buffer when the MDB is stopped by the administrator.
receiveCCSID	int	<ul style="list-style-type: none"> <li><b>0</b> - Use JVM Charset.defaultCharset</li> <li>1208 - UTF-8</li> <li>A supported coded character set identifier</li> </ul>	Destination property that sets the target CCSID for queue manager message conversion. The value is ignored unless <b>receiveConversion</b> is set to QMGR.
receiveConversion	String	<ul style="list-style-type: none"> <li><b>CLIENT_MSG</b></li> <li>QMGR</li> </ul>	Destination property that determines whether data conversion is going to be performed by the queue manager.
sharedSubscription	Boolean	<ul style="list-style-type: none"> <li><b>False</b> - The MDB should not open the subscription as a shared subscription.</li> <li>True - The MDB should open the subscription as a shared subscription (with the rules that JMS 2.0 implies, see the JMS 2.0 specification at Java.net).</li> </ul>	Controls how an MDB is driven from a shared subscription. For more information about how to use this property, see "Examples of how to define the sharedSubscription property" on page 1120.

Table 160. Properties of an ActivationSpec object that are used to create a JMS connection consumer (continued)

Name of property	Type	Valid values (default value in bold)	Description
startTimeout	int	<ul style="list-style-type: none"> <li>• <b>10 000</b></li> <li>• A positive integer</li> </ul>	The time, in milliseconds, within which delivery of a message to an MDB must start after the work to deliver the message has been scheduled. If this time elapses, the message is rolled back onto the queue.
subscriptionDurability	String	<ul style="list-style-type: none"> <li>• <b>NonDurable</b> - A nondurable subscription is used to deliver messages to an MDB subscribing to the topic.</li> <li>• <b>Durable</b> - A durable subscription is used to deliver messages to an MDB subscribing to the topic.</li> </ul>	Whether a durable or nondurable subscription is used to deliver messages to an MDB subscribing to the topic
subscriptionName	String	<ul style="list-style-type: none"> <li>• <b>"" (empty string)</b></li> <li>• A subscription name</li> </ul>	The name of the durable subscription
useJNDI	Boolean	<ul style="list-style-type: none"> <li>• <b>false</b> - The property called destination is interpreted as the name of an IBM MQ queue or a topic.</li> <li>• <b>true</b> - The property called destination is interpreted as the name of a javax.jms.Queue object or javax.jms.Topic object in the JNDI namespace of the application server.</li> </ul>	Determines how the value of the property called destination is interpreted <b>Note:</b> This property is deprecated in IBM MQ Version 9.0. The destinationLookup property should be used instead.

### Property conflicts and dependencies

An ActivationSpec object can have conflicting properties. For example, you can specify TLS properties for a connection in bindings mode. In this case, the behavior is determined by the transport type and the messaging domain, which is either point-to-point or publish/subscribe as determined by the **destinationType** property. Any properties that are not applicable to the specified transport type or messaging domain are ignored.

If you define a property that requires other properties to be defined, but you do not define these other properties, the ActivationSpec object throws an InvalidPropertyException exception when its validate() method is called during the deployment of an MDB. The exception is reported to the administrator of the application server in a manner that depends on the application server. For example, if you set the subscriptionDurability property to Durable, indicating that you want use durable subscriptions, you must also define the **subscriptionName** property.

If the properties called **ccdtURL** and **channel** are both defined, an InvalidPropertyException exception is thrown. However, if you define the **ccdtURL** property only, leaving the property called **channel** with its default value of SYSTEM.DEF.SVRCONN, no exception is thrown, and the client channel definition table identified by the **ccdtURL** property is used to start a JMS connection.

### ActivationSpec connectionFactoryLookup and destinationLookup properties

These two properties can be used to specify the JNDI names of ConnectionFactory and Destination objects that are used in preference to the properties of the ActivationSpec as defined in Table 159 on page 1107 and Table 160 on page 1115.

It is important to note the following points that describe how these properties work in detail.

### **connectionFactoryLookup**

The `ConnectionFactory` that is looked up from JNDI is used as a source of the properties listed in Table 159 on page 1107. The `ConnectionFactory` object is not used to actually create any JMS connections, only the properties of the object are queried. These properties from the `ConnectionFactory` override any properties that are defined on the `ActivationSpec`. There is a single exception to this. If the `ActivationSpec` has the `ClientID` property set, then the value of this property overrides the value specified in the `ConnectionFactory`. This is because a common scenario is using a single `ConnectionFactory` with multiple `ActivationSpecs`. This simplifies administration. However, the JMS 2.0 specification states that every JMS Connection created from a `ConnectionFactory` should have a unique `ClientID`. Because of this, `ActivationSpecs` need to have the ability to override any value set on the `ConnectionFactory`. If no `ClientID` is set on the `ActivationSpec`, any value on the connection factory is used.

### **destinationLookup**

A `Destination` and a `UseJndi` property are defined on the `ActivationSpec`. If the `UseJndi` flag is set to true, then the text specified in the destination property is considered to be a JNDI name and a destination object with that JNDI name is looked up from JNDI.

The `destinationLookup` property behaves in exactly the same way. If it has been set, then a destination object with the JNDI name specified by the property is looked up from JNDI. This property has precedence over the `useJNDI` property.

The `useJNDI` property is deprecated at IBM MQ Version 9.0 as the `destinationLookup` property is the JMS 2.0 specification equivalent of performing the same function.

### **ActivationSpec properties with no equivalents in IBM MQ classes for JMS**

Most of the properties of an `ActivationSpec` object are equivalent to properties of IBM MQ classes for JMS objects or parameters of IBM MQ classes for JMS methods. However, three tuning properties, and one usability property, have no equivalents in IBM MQ classes for JMS:

#### **startTimeout**

The time, in milliseconds, that the work manager of the application server waits for resources to become available after the resource adapter schedules a `Work` object to deliver a message to an MDB. If this time elapses before delivery of the message starts, the `Work` object times out, the message is rolled back onto the queue, and the resource adapter can then attempt to deliver the message again. A warning is written to diagnostic trace, if enabled, but does not otherwise affect the process of delivering messages. You might expect this condition to occur only at times when the application server is experiencing a very high load. If the condition occurs regularly, consider increasing the value of this property to give the work manager longer to schedule message delivery.

#### **maxPoolDepth**

The maximum number of server sessions in the server session pool used by a connection consumer. When a server session is created, it starts a conversation with a queue manager. The connection consumer uses a server session to deliver a message to an MDB. A larger pool depth allows more messages to be delivered concurrently in high volume situations, but uses more resources of the application server. If many MDBs are to be deployed, consider making the pool depth smaller in order to maintain the load on the application server at a manageable level. Each connection consumer uses its own server session pool, so that this property does not define the total number of server sessions available to all connection consumers.

#### **poolTimeout**

The time, in milliseconds, that an unused server session is held open in the server session pool before being closed due to inactivity. A transient increase in the message workload causes additional server sessions to be created in order to distribute the load but, after the message workload returns to normal, the additional server sessions remain in the pool and are not used.

Every time a server session is used, it is marked with a timestamp. Periodically a scavenger thread checks that each server session has been used within the period specified by this property. If a server session has not been used, it is closed and removed from the server session pool. A server session might not be closed immediately after the specified period has elapsed, this property represents the minimum period of inactivity before removal.

#### **useJNDI**

For a description of this property, see Table 160 on page 1115.

### **Deploying an MDB**

To deploy an MDB, first define the properties of an `ActivationSpec` object, specifying the properties that the MDB requires. The following example is a typical set of properties that you might define explicitly:

```
channel: SYSTEM.DEF.SVRCONN
destination: SYSTEM.DEFAULT.LOCAL.QUEUE
destinationType: javax.jms.Queue
hostName: 192.168.0.42
messageSelector: color='red'
port: 1414
queueManager: ExampleQM
transportType: CLIENT
```

The application server uses the properties to create an `ActivationSpec` object, which is then associated with an MDB. The properties of the `ActivationSpec` object determine how messages are delivered to the MDB. Deployment of the MDB fails if the MDB requires distributed transactions but the resource adapter does not support distributed transactions. For information about how to install the resource adapter so that distributed transactions are supported, see “Installing the IBM MQ resource adapter” on page 1099.

If more than one MDB is receiving messages from the same destination, then a message sent in the point-to-point domain is received by only one MDB, even if other MDBs are eligible to receive the message. In particular, if two MDBs are using different message selectors, and an incoming message matches both message selectors, only one of the MDBs receives the message. The MDB chosen to receive a message is undefined, and you cannot rely on a specific MDB receiving the message. Messages sent in the publish/subscribe domain are received by all eligible MDBs.

In some circumstances, a message delivered to an MDB might be rolled back onto an IBM MQ queue. This roll back can happen, for example, if a message is delivered within a unit of work that is then rolled back. A message that is rolled back is delivered again, but a badly formatted message might repeatedly cause an MDB to fail and therefore cannot be delivered. Such a message is called a poison message. You can configure IBM MQ so that IBM MQ classes for JMS automatically transfers a poison message to another queue for further investigation or discards the message.

For details on how to handle poison messages, see “Handling poison messages in IBM MQ classes for JMS” on page 956.

**Related information:**

Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client  
 Federal Information Processing Standards (FIPS) for UNIX, Linux and Windows  
 Configuring JMS resources in WebSphere Application Server

*Examples of how to define the sharedSubscription property:*

You can define the sharedSubscription property of an activation specification within a WebSphere Application Server Liberty server.xml file. Alternatively, you can define the property within a message driven bean (MDB) using annotations.

**Example: defining within a Liberty server.xml file**

Within a WebSphere Application Server Liberty server.xml file, you define an activation specification as shown in the following example. This example creates a durable shared subscription to a queue manager on localhost/port 1490.

```
<jmsActivationSpec id="SubApp/SubscribingEJB/SubscribingMDB" authDataRef="JMSConnectionAlias">
<properties.wmqJms hostName="localhost" port="1490" maxPoolDepth="5" subscriptionName="MySubName"
subscriptionDurability="DURABLE" sharedSubscription="true"/>
</jmsActivationSpec>
```

**Example: defining within an MDB**

You can also define the sharedSubscription property within the MDB using annotations as shown in the following example:

```
@ActioncationConfigProperty(propertyName = "sharedSubscription",
propertyValue = "true")
```

The following example shows a piece of MDB code that uses the annotations method:

```
/**
 * Message-Driven Bean example using Annotations for configuration
 */
@MessageDriven(
 activationConfig = {
 @ActivationConfigProperty(
 propertyName = "destinationType", propertyValue = "javax.jms.Topic"),
 @ActivationConfigProperty(
 propertyName = "sharedSubscription", propertyValue = "TRUE"),
 @ActivationConfigProperty(
 propertyName = "destination", propertyValue = "JNDI_TOPIC_NAME")
 },
 mappedName = "Stock/IBM")
public class SubscribingMDB implements MessageListener {

 // Default constructor.
 public SubscribingMDB() {
 }

 // @see MessageListener#onMessage(Message)
 public void onMessage(Message message) {
 // implement business logic here
 }
}
```



**Related concepts:**

“Cloned and shared subscriptions” on page 999

In IBM MQ Version 8.0 or later, there are two methods for giving multiple consumers access to the same subscription. These two methods are by using cloned subscriptions, or by using shared subscriptions.

**Related information:**

Subscribers and subscriptions

Subscription durability

**Configuring the resource adapter for outbound communication:**

To configure outbound communication, define the properties of a `ConnectionFactory` object and an administered destination object.

**Example of using outbound communication**

When using outbound communication, an application running in an application server starts a connection to a queue manager, and then sends messages to its queues and receives messages from its queues in a synchronous manner. For example, the following servlet method, `doGet()`, uses outbound communication:

```
protected void doGet(HttpServletRequest request, HttpServletResponse response)
 throws ServletException, IOException {
 ...

 // Look up ConnectionFactory and Queue objects from the JNDI namespace
 InitialContext ic = new InitialContext();
 ConnectionFactory cf = (javax.jms.ConnectionFactory) ic.lookup("myCF");
 Queue q = (javax.jms.Queue) ic.lookup("myQueue");

 // Create and start a connection
 Connection c = cf.createConnection();
 c.start();

 // Create a session and message producer
 Session s = c.createSession(false, Session.AUTO_ACKNOWLEDGE);
 MessageProducer pr = s.createProducer(q);

 // Create and send a message
 Message m = s.createTextMessage("Hello, World!");
 pr.send(m);

 // Create a message consumer and receive the message just sent
 MessageConsumer co = s.createConsumer(q);
 Message mr = co.receive(5000);

 // Close the connection
 c.close();
}
```

When the servlet receives an HTTP GET request, it retrieves a `ConnectionFactory` object and a `Queue` object from the JNDI namespace, and uses the objects to send a message to an IBM MQ queue. The servlet then receives the message that it has sent.

## Resources needed for outbound communication

To configure outbound communication, define Java EE Connector Architecture (JCA) resources in the following categories:

- The properties of a ConnectionFactory object, which the application server uses to create a JMS ConnectionFactory object.
- The properties of an administered destination object, which the application server uses to create a JMS Queue object or JMS Topic object.

The way you define these properties depends on the administration interfaces provided by your application server. ConnectionFactory, Queue, and Topic objects created by the application server are bound into a JNDI namespace from where they can be retrieved by an application.

Typically, you define one ConnectionFactory object for each queue manager that applications might need to connect to. You define one Queue object for each queue that applications might need to access in the point-to-point domain. And you define one Topic object for each topic that applications might want to publish or subscribe to. A ConnectionFactory object can be domain independent. Alternatively, it can be domain-specific, a QueueConnectionFactory object for the point-to-point domain or a TopicConnectionFactory object for the publish/subscribe domain.

**Tip:** With JMS 2.0, a connection factory can be used to create both connections and contexts. As a result, it is possible to have a connection pool associated with a connection factory that contains a mixture of both connections and contexts. It is recommended that a connection factory is only used for creating connections or creating contexts. This ensures that the connection pool for that connection factory only contains objects of a single type, which makes the pool more efficient.

## Properties of a ConnectionFactory object

Table 161 lists the properties of a ConnectionFactory object. The application server uses these properties to create a JMS ConnectionFactory object.

Table 161. Properties of a ConnectionFactory object

Name of property	Type	Valid values (default value in bold)	Description
applicationName	String	<ul style="list-style-type: none"> <li>• The invoking class name, if it is available, adjusted to be no longer than 28 characters. If it is not available, the string WebSphere MQ Client for Java is used.</li> </ul>	The name by which an application is registered with the queue manager. This application name is shown by the <b>DISPLAY CONN MQSC/PCF</b> command (where the field is called <b>APPLTAG</b> ) or in the IBM MQ Explorer Application Connections display (where the field is called <b>App name</b> ).
brokerCCSubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.JMS.ND.CC.SUBSCRIBER.QUEUE</b></li> <li>• A queue name</li> </ul>	The name of the queue from which a connection consumer receives nondurable subscription messages.
brokerControlQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.BROKER.CONTROL.QUEUE</b></li> <li>• A queue name</li> </ul>	The name of the broker control queue.
brokerPubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.BROKER.DEFAULT.STREAM</b></li> <li>• A queue name</li> </ul>	The name of the queue where published messages are sent (the stream queue).
brokerQueueManager <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• "" (<b>empty string</b>)</li> <li>• A queue manager name</li> </ul>	The name of the queue manager on which the broker is running.

Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
brokerSubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.JMS.ND.SUBSCRIBER.QUEUE</b></li> <li>• A queue name</li> </ul>	<p>The name of the queue from which a nondurable message consumer receives messages.</p> <p>See BROKERSUBQ property for more information.</p>
brokerVersion <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>unspecified</b> - After the broker has been migrated from V6 to V7, set this property so that RFH2 headers are no longer used. After migration this property is no longer relevant.</li> <li>• <b>V1</b> - To use an IBM MQ Publish/Subscribe broker. This value is the default value if TRANSPORT is set to BIND or CLIENT.</li> <li>• <b>V2</b> - To use a broker of IBM Integration Bus in native mode. This value is the default value if TRANSPORT is set to DIRECT or DIRECTHTTP.</li> </ul>	The version of the broker being used.
ccdtURL	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A uniform resource locator (URL)</li> </ul>	A URL that identifies the name and location of the file containing the client channel definition table (CCDT) and specifies how the file can be accessed.
CCSID	String	<ul style="list-style-type: none"> <li>• <b>819</b></li> <li>• A coded character set identifier supported by the Java virtual machine (JVM)</li> </ul>	The coded character set identifier for a connection.
channel	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.DEF.SVRCONN</b></li> <li>• The name of an MQI channel</li> </ul>	The name of the MQI channel to use.
cleanupInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>3 600 000</b></li> <li>• A positive integer</li> </ul>	The interval, in milliseconds, between background runs of the publish/subscribe cleanup utility.
cleanupLevel <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>SAFE</b></li> <li>• NONE</li> <li>• STRONG</li> <li>• FORCE</li> <li>• NONDUR</li> </ul>	The cleanup level for a broker-based subscription store.
clientID	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A client identifier</li> </ul>	The client identifier for a connection.
cloneSupport	String	<ul style="list-style-type: none"> <li>• <b>DISABLED</b> - Only one instance of a durable topic subscriber can run at a time.</li> <li>• <b>ENABLED</b> - Two or more instances of the same durable topic subscriber can run simultaneously, but each instance must run in a separate Java virtual machine (JVM).</li> </ul>	Whether two or more instances of the same durable topic subscriber can run simultaneously.

Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
connectionNameList	String	<ul style="list-style-type: none"> <li>• <b>localhost(1414)</b></li> <li>• A string composed of items separated by commas where each item takes the format: <i>HOSTNAME (PORT)</i></li> </ul> <p>where <i>HOSTNAME</i> is either a DNS name or an IP address.</p>	<p>A list of TCP/IP connection names used for outbound communications.</p> <p><b>connectionNameList</b> supersedes the <b>hostname</b> and <b>port</b> properties.</p> <p>This property is used to reconnect to multi-instance queue managers.</p> <p><b>connectionNameList</b> is similar in form to <b>localAddress</b>, but must not be confused with it. <b>LocalAddress</b> specifies the characteristics of the local communications, whereas <b>connectionNameList</b> specifies how to reach a remote queue manager.</p>
failIfQuiesce	Boolean	<ul style="list-style-type: none"> <li>• <b>true</b></li> <li>• false</li> </ul>	Whether calls to certain methods fail if the queue manager is in a quiescing state.
headerCompression	String	<ul style="list-style-type: none"> <li>• <b>NONE</b></li> <li>• SYSTEM - RLE message header compression is performed.</li> </ul>	A list of the techniques that can be used for compressing header data on a connection.
hostName	String	<ul style="list-style-type: none"> <li>• <b>localhost</b></li> <li>• A host name</li> <li>• An IP address</li> </ul>	<p>The host name or IP address of the system on which the queue manager resides.</p> <p>The <b>hostname</b> and <b>port</b> properties are superseded by the <b>connectionNameList</b> property when it is specified.</p>
localAddress	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string in the format: [ <i>host_name</i> ][(<i>low_port</i> [, <i>high_port</i> ])]</li> </ul> <p>where <i>host_name</i> is a host name or IP address, <i>low_port</i> and <i>high_port</i> are TCP port numbers, and brackets denote an optional component</p>	<p>For a connection to a queue manager, this property specifies either or both of the following:</p> <ul style="list-style-type: none"> <li>• The local network interface to be used</li> <li>• The local port, or range of local ports, to be used</li> </ul> <p><b>LocalAddress</b> is similar in form to <b>connectionNameList</b>, but must not be confused with it. <b>LocalAddress</b> specifies the characteristics of the local communications, whereas <b>connectionNameList</b> specifies how to reach a remote queue manager.</p>

Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
messageCompression	String	<ul style="list-style-type: none"> <li>• <b>NONE</b></li> <li>• A list of one or more of the following values separated by blank characters: RLE ZLIBFAST ZLIBHIGH</li> </ul>	A list of the techniques that can be used for compressing message data on a connection.
messageSelection <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>CLIENT</b></li> <li>• BROKER</li> </ul>	Determines whether message selection is done by IBM MQ classes for JMS or by the broker. Message selection by the broker is not supported when <b>brokerVersion</b> has the value 1.
password	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A password</li> </ul>	The default password to use when creating a connection to the queue manager.
pollingInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>5000</b></li> <li>• Any positive integer</li> </ul>	If each message listener within a session has no suitable message on its queue, this value is the maximum interval, in milliseconds, that elapses before each message listener tries again to get a message from its queue. If it frequently happens that no suitable message is available for any of the message listeners in a session, consider increasing the value of this property. This property is relevant only if <b>TRANSPORT</b> has the value BIND or CLIENT.
port	int	<ul style="list-style-type: none"> <li>• <b>1414</b></li> <li>• A TCP port number</li> </ul>	<p>The port on which the queue manager listens.</p> <p>The <b>hostname</b> and <b>port</b> properties are superseded by the <b>connectionNameList</b> property when it is specified.</p>
providerVersion	string	<ul style="list-style-type: none"> <li>• <b>unspecified</b></li> <li>• A string in one of the following formats <ul style="list-style-type: none"> <li>– V.R.M.F</li> <li>– V.R.M</li> <li>– V.R</li> <li>– V</li> </ul> </li> </ul> <p>where V, R, M, and F are integer values greater than or equal to zero.</p>	The version, release, modification level and fix pack of the queue manager to which the application intends to connect.
pubAckInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>25</b></li> <li>• A positive integer</li> </ul>	The number of messages published by a publisher before IBM MQ classes for JMS requests an acknowledgment from the broker.
queueManager	String	<ul style="list-style-type: none"> <li>• <b>"" (empty string)</b></li> <li>• A queue manager name</li> </ul>	The name of the queue manager to connect to.

Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
receiveExit <sup>3</sup>	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string comprising one or more items separated by commas, where each item is the fully qualified name of a class that implements the IBM MQ classes for Java interface, MQReceiveExit</li> </ul>	Identifies a channel receive exit program, or a sequence of receive exit programs to be run in succession.
receiveExitInit	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string comprising one or more items of user data separated by commas</li> </ul>	The user data that is passed to channel receive exit programs when they are called.
rescanInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>5000</b></li> <li>• Any positive integer</li> </ul>	When a message consumer in the point-to-point domain uses a message selector to select which messages it wants to receive, IBM MQ classes for JMS searches the IBM MQ queue for suitable messages in the sequence determined by the <b>MsgDeliverySequence</b> attribute of the queue. When IBM MQ classes for JMS finds a suitable message and delivers it to the consumer, IBM MQ classes for JMS resumes the search for the next suitable message from its current position in the queue. IBM MQ classes for JMS continues to search the queue in this way until it reaches the end of the queue, or until the interval of time in milliseconds, as determined by the value of this property, has expired. In each case, IBM MQ classes for JMS returns to the beginning of the queue to continue its search, and a new time interval commences.
securityExit <sup>3</sup>	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• The fully qualified name of a class that implements the IBM MQ classes for Java interface, MQSecurityExit</li> </ul>	Identifies a channel security exit program.
securityExitInit	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string of user data</li> </ul>	The user data that is passed to a channel security exit program when it is called.
sendCheckCount	int	<ul style="list-style-type: none"> <li>• <b>0</b></li> <li>• Any positive integer</li> </ul>	The number of send calls to allow between checking for asynchronous put errors, within a single non-transacted JMS session.
sendExit <sup>3</sup>	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string comprising one or more items separated by commas, where each item is the fully qualified name of a class that implements the IBM MQ classes for Java interface, MQSendExit</li> </ul>	Identifies a channel send exit program, or a sequence of send exit programs to be run in succession.

Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
sendExitInit	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string comprising one or more items of user data separated by commas</li> </ul>	The user data that is passed to channel send exit programs when they are called.
shareConvAllowed	Boolean	<ul style="list-style-type: none"> <li>• <b>NO</b> - A client connection cannot share its socket.</li> <li>• <b>YES</b> - A client connection can share its socket.</li> </ul>	Whether a client connection can share its socket with other top-level JMS connections from the same process to the same queue manager, if the channel definitions match.
sparseSubscriptions <sup>1</sup>	Boolean	<ul style="list-style-type: none"> <li>• <b>false</b> - Subscriptions receive frequent matching messages.</li> <li>• <b>true</b> - Subscriptions receive infrequent matching messages. This value requires that the subscription queue can be opened for browse.</li> </ul>	Controls the message retrieval policy of a TopicSubscriber object.
sslCertStores	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A string of one or more LDAP URLs separated by blanks. Each LDAP URL has the format:  <code>ldap://host_name [: port ]</code>                       where <i>host_name</i> is a host name or IP address, <i>port</i> is a TCP port number, and brackets denote an optional component.</li> </ul>	The Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs) for use on a TLS connection.
sslCipherSuite	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• The name of a CipherSuite</li> </ul>	The CipherSuite to use for a TLS connection.
sslFipsRequired <sup>2</sup>	Boolean	<ul style="list-style-type: none"> <li>• <b>false</b></li> <li>• <b>true</b></li> </ul>	Whether a TLS connection must use a CipherSuite that is supported by the IBM Java JSSE FIPS provider (IBMJSSEFIPS).
sslPeerName	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A template for distinguished names</li> </ul>	For a TLS connection, a template that is used to check the distinguished name in the digital certificate provided by the queue manager.
sslResetCount	int	<ul style="list-style-type: none"> <li>• <b>0</b></li> <li>• An integer in the range 0 - 999 999 999</li> </ul>	The total number bytes sent and received by a TLS connection before the secret keys used by TLS are renegotiated.
sslSocketFactory	String	A string representing the fully qualified class name of a class providing an implementation of the javax.net.ssl.SSLSocketFactory interface, optionally including an argument to be passed to the constructor method, enclosed in parentheses.	Any connections established in the scope of the administered destination object use sockets obtained from this implementation of the SSLSocketFactory interface.
statusRefreshInterval <sup>1</sup>	int	<ul style="list-style-type: none"> <li>• <b>60000</b></li> <li>• Any positive integer</li> </ul>	The interval, in milliseconds, between refreshes of the long running transaction that detects when a subscriber loses its connection to the queue manager. This property is relevant only if <b>SUBSTORE</b> has the value <b>QUEUE</b> .

Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
subscriptionStore <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>BROKER</b></li> <li>• MIGRATE</li> <li>• QUEUE</li> </ul>	Determines where IBM MQ classes for JMS stores persistent data about active subscriptions.
targetClientMatching	Boolean	<ul style="list-style-type: none"> <li>• <b>true</b></li> <li>• false</li> </ul>	Whether a reply message, sent to the queue identified by the JMSReplyTo header field of an incoming message, has an MQRFH2 header only if the incoming message has an MQRFH2 header.



Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
temporaryModel	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.DEFAULT.MODEL.QUEUE</b></li> <li>• SYSTEM.JMS.TEMPQ.MODEL</li> <li>• Any string</li> </ul>	<p>The name of the model queue from which JMS temporary queues are created.</p> <p>Use <b>SYSTEM.DEFAULT.MODEL.QUEUE</b> if both of the following statements are true:</p> <ul style="list-style-type: none"> <li>• Your application uses a temporary queue that will accept non-persistent messages.</li> <li>• Only one application will create a temporary queue on the queue manager that the ConnectionFactory points to at a time. Note that <b>SYSTEM.DEFAULT.MODEL.QUEUE</b> can only be opened by one application at a time.</li> </ul> <p>Use <b>SYSTEM.JMS.TEMPQ.MODEL</b> in the following situations:</p> <ul style="list-style-type: none"> <li>• When your application uses a temporary queue that will accept persistent messages.</li> <li>• If multiple applications can connect to the queue manager that the ConnectionFactory points to and those applications need to create temporary queues at the same time.</li> </ul> <p>Define a new model queue with the <b>DEFPSIST</b> attribute set to YES, and the <b>DEFSOPT</b> attribute set to SHARED in the following situation:</p> <ul style="list-style-type: none"> <li>• When your application uses a temporary queue that will accept non-persistent messages, and multiple applications will connect to the queue manager that the ConnectionFactory points to, and those applications need to create temporary queues at the same time.</li> </ul> <p>When the new model queue is created, set the <b>temporaryModel</b> property to the name of the new model queue.</p>

Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
tempQPrefix	String	<ul style="list-style-type: none"> <li>• "" (<b>empty string</b>)</li> <li>• A prefix that can be used to form the name of an IBM MQ dynamic queue. The rules for forming the prefix are the same as the rules for forming the contents of the <b>DynamicQName</b> field in an IBM MQ object descriptor, structure MQOD, but the last non-blank character must be an asterisk (*). If the value of the property is the empty string, IBM MQ classes for JMS uses the value AMQ.* when creating a dynamic queue.</li> </ul>	The prefix that is used to form the name of an IBM MQ dynamic queue.
tempTopicPrefix	String	Any non-null string consisting only of valid characters for an IBM MQ topic string	When creating temporary topics, JMS generates a topic string of the form "TEMP/ <i>TEMPTOPICPREFIX/unique_id</i> ", or if this property is left with the default value, just "TEMP/ <i>unique_id</i> ". Specifying a non-empty <b>TEMPTOPICPREFIX</b> allows specific model queues to be defined for creating the managed queues for subscribers to temporary topics created under this connection.
transportType	String	<ul style="list-style-type: none"> <li>• <b>CLIENT</b></li> <li>• <b>BINDINGS</b></li> <li>• <b>BINDINGS_THEN_CLIENT</b></li> </ul>	<p>Whether a connection to a queue manager uses client mode or bindings mode. If the value <b>BINDINGS_THEN_CLIENT</b> is specified, the resource adapter first tries to make a connection in bindings mode. If this connection attempt fails, the resource adapter then tries to make a client mode connection.</p> <p><b>z/OS</b> If an activation specification that is running on a WebSphere Application Server for z/OS system has been configured to use the <b>BINDINGS_THEN_CLIENT</b> transport mode and a previously established connection is broken, then any reconnection attempts by the activation specification first attempt to use the <b>BINDINGS</b> transport mode. If the <b>BINDINGS</b> transport mode connection attempt is unsuccessful, the activation specification subsequently attempts a <b>CLIENT</b> transport mode connection.</p>
username	String	<ul style="list-style-type: none"> <li>• <b>null</b></li> <li>• A user name</li> </ul>	The default user name to use when creating a connection to a queue manager.

Table 161. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
wildcardFormat	int	<ul style="list-style-type: none"> <li>• <b>CHAR</b>- Recognizes character wildcards only, as used in broker version 1</li> <li>• <b>TOPIC</b> - Recognizes topic level wildcards only, as used in broker version 2</li> </ul>	Which version of wildcard syntax is to be used.

**Notes:**

1. This property can be used with Version 7.0 of IBM WebSphere MQ classes for JMS but does not affect an application connected to a Version 7.0 queue manager unless the providerVersion property is set to a version number less than 7.
2. For important information about using the sslFipsRequired property, see “Limitations of the IBM MQ resource adapter” on page 1096.
3. For information on how to configure the resource adapter so that it can locate an exit, see “Configuring IBM MQ classes for JMS to use channel exits” on page 979.

The following example shows a typical set of properties of a ConnectionFactory object:

```
channel: SYSTEM.DEF.SVRCONN
hostName: 192.168.0.42
port: 1414
queueManager: ExampleQM
transportType: CLIENT
```

**Properties of an administered destination object**

The application server uses the properties of an administered destination object to create a JMS Queue object or JMS Topic object.

Table 162 lists the properties that are common to a Queue object and a Topic object.

Table 162. Properties that are common to a Queue object and a Topic object

Name of property	Type	Valid values (default value in bold)	Description
CCSID	String	<ul style="list-style-type: none"> <li>• <b>1208</b></li> <li>• A coded character set identifier supported by the Java virtual machine (JVM)</li> </ul>	The coded character set identifier for the destination.
encoding	String	<ul style="list-style-type: none"> <li>• <b>NATIVE</b></li> <li>• A string of three characters: <ul style="list-style-type: none"> <li>- The first character specifies the representation of binary integers: <ul style="list-style-type: none"> <li>- <i>N</i> denotes normal encoding.</li> <li>- <i>R</i> denotes reverse encoding.</li> </ul> </li> <li>- The second character specifies the representation of packed decimal integers: <ul style="list-style-type: none"> <li>- <i>N</i> denotes normal encoding.</li> <li>- <i>R</i> denotes reverse encoding.</li> </ul> </li> <li>- The third character specifies the representation of floating point numbers: <ul style="list-style-type: none"> <li>- <i>N</i> denotes standard IEEE encoding.</li> <li>- <i>R</i> denotes reverse IEEE encoding.</li> <li>- <i>3</i> denotes zSeries encoding.</li> </ul> </li> </ul> </li> </ul> <p>NATIVE is equivalent to the string NNN.</p>	The representation of binary integers, packed decimal integers, and floating point numbers for the destination.

Table 162. Properties that are common to a Queue object and a Topic object (continued)

Name of property	Type	Valid values (default value in bold)	Description
expiry	String	<ul style="list-style-type: none"> <li>• <b>APP</b> - The expiry time of a message is determined by the message producer.</li> <li>• UNLIM - A message never expires.</li> <li>• 0 - A message never expires.</li> <li>• A positive integer representing the expiry time of a message in milliseconds.</li> </ul>	The expiry time of a message sent to the destination.
failIfQuiesce	String	<ul style="list-style-type: none"> <li>• <b>true</b></li> <li>• false</li> </ul>	Whether an attempt to access the destination fails if the queue manager is in a quiescing state.
messageBodyStyle	String	<ul style="list-style-type: none"> <li>• <b>UNSPECIFIED</b></li> <li>• JMS</li> <li>• MQ</li> </ul>	<p>You can set the <b>messageBodyStyle</b> property on JMS queues and topics: UNSPECIFIED(default)</p> <ul style="list-style-type: none"> <li>• When sending, IBM MQ classes for JMS generate and include an MQRFH2 header, depending on the value of WMQ_TARGET_CLIENT.</li> <li>• When receiving, IBM MQ classes for JMS set the JMS message properties according to values in the MQRFH2, if present. MQRFH2 is not presented as part of the JMS message body.</li> </ul> <p>JMS</p> <ul style="list-style-type: none"> <li>• When sending, IBM MQ classes for JMS automatically generates an MQRFH2 header and includes the header in the IBM MQ message.</li> <li>• When receiving, IBM MQ classes for JMS set the JMS message properties according to values in the MQRFH2, if present. MQRFH2 is not presented as part of the JMS message body.</li> </ul> <p>MQ</p> <ul style="list-style-type: none"> <li>• When sending, IBM MQ classes for JMS do not generate an MQRFH2.</li> <li>• When receiving, IBM MQ classes for JMS present the MQRFH2 as part of the JMS message body.</li> </ul>

Table 162. Properties that are common to a Queue object and a Topic object (continued)

Name of property	Type	Valid values (default value in bold)	Description
persistence	String	<ul style="list-style-type: none"> <li>• <b>APP</b> - The persistence of a message is determined by the message producer.</li> <li>• <b>QDEF</b> - The persistence of a message is determined by the <b>DefPersistence</b> attribute of the IBM MQ queue.</li> <li>• <b>PERS</b> - A message is persistent.</li> <li>• <b>NON</b> - A message is nonpersistent.</li> <li>• <b>HIGH</b> - The persistence of a message is determined by the <b>NonPersistentMessageClass</b> attribute of the IBM MQ queue according to the explanation in “JMS persistent messages” on page 970.</li> </ul>	The persistence of a message sent to the destination.
priority	String	<ul style="list-style-type: none"> <li>• <b>APP</b> - The priority of a message is determined by the message producer.</li> <li>• <b>QDEF</b> - The priority of a message is determined by the <b>DefPriority</b> attribute of the IBM MQ queue.</li> <li>• An integer in the range 0, lowest priority, to 9, highest priority.</li> </ul>	The priority of a message sent to the destination.
putAsyncAllowed	String	<ul style="list-style-type: none"> <li>• <b>QUEUE</b> - Determine whether asynchronous puts are allowed by referring to the queue definition.</li> <li>• <b>TOPIC</b> - Determine whether asynchronous puts are allowed by referring to the topic definition.</li> <li>• <b>DESTINATION</b> - Determine whether asynchronous puts are allowed by referring to the queue or topic definition.</li> <li>• <b>DISABLED</b> - Asynchronous puts are not allowed.</li> <li>• <b>ENABLED</b> - Asynchronous puts are allowed.</li> </ul>	Whether message producers are allowed to use asynchronous puts to send messages to this destination.
readAheadAllowed	int	<ul style="list-style-type: none"> <li>• <b>DESTINATION</b> - Determine whether read ahead is allowed by referring to the queue or topic definition.</li> <li>• <b>DISABLED</b> - Read ahead is not allowed.</li> <li>• <b>ENABLED</b> - Read ahead is allowed.</li> <li>• <b>QUEUE</b> - Determine whether read ahead is allowed by referring to the queue definition.</li> <li>• <b>TOPIC</b> - Determine whether read ahead is allowed by referring to the topic definition.</li> </ul>	Whether message consumers and queue browsers are allowed to use read ahead to get nonpersistent messages from the destination into an internal buffer before receiving them.
receiveCCSID	int	<ul style="list-style-type: none"> <li>• <b>0</b> - Use JVM <code>Charset.defaultCharset</code></li> <li>• 1208 - UTF-8</li> <li>• A supported coded character set identifier</li> </ul>	Destination property that sets the target CCSID for queue manager message conversion. The value is ignored unless <b>receiveConversion</b> is set to QMGR.
receiveConversion	String	<ul style="list-style-type: none"> <li>• <b>CLIENT_MSG</b></li> <li>• <b>QMGR</b></li> </ul>	Destination property that determines whether data conversion is going to be performed by the queue manager.

Table 162. Properties that are common to a Queue object and a Topic object (continued)

Name of property	Type	Valid values (default value in bold)	Description
targetClient	String	<ul style="list-style-type: none"> <li>• <b>JMS</b> - The target of a message is a JMS application.</li> <li>• MQ - The target of a message is a non-JMS IBM MQ application.</li> </ul>	Whether the target of a message sent to the destination is a JMS application. A message with a target that is a JMS application contains an MQRFH2 header.

Table 163 lists the properties that are specific to a Queue object.

Table 163. Properties that are specific to a Queue object

Name of property	Type	Valid values (default value in bold)	Description
baseQueueManagerName	String	<ul style="list-style-type: none"> <li>• "" (<b>empty string</b>)</li> <li>• A queue manager name</li> </ul>	The name of the queue manager that owns the underlying IBM MQ queue.
baseQueueName	String	<ul style="list-style-type: none"> <li>• "" (<b>empty string</b>)</li> <li>• A queue name</li> </ul>	The name of the underlying IBM MQ queue.

Table 164 lists the properties that are specific to a Topic object.

Table 164. Properties that are specific to a Topic object

Name of property	Type	Valid values (default value in bold)	Description
baseTopicName	String	<ul style="list-style-type: none"> <li>• "" (<b>empty string</b>)</li> <li>• A topic name</li> </ul>	The name of the underlying topic.
brokerCCDurSubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.JMS.D.CC.SUBSCRIBER.QUEUE</b></li> <li>• A queue name</li> </ul>	The name of the queue from which a connection consumer receives durable subscription messages.
brokerDurSubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>SYSTEM.JMS.D.SUBSCRIBER.QUEUE</b></li> <li>• A queue name</li> </ul>	The name of the queue from which a durable topic subscriber receives messages. See the <b>BROKEDURRSUBQ</b> property in the IBM MQ Explorer documentation for more information.
brokerPubQueue <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>Not set</b></li> <li>• A queue name</li> </ul>	The name of the queue where published messages are sent (the stream queue). The value of this property overrides the value of the <b>brokerPubQueue</b> property of the ConnectionFactory object. However, if you do not set the value of this property, the value of the <b>brokerPubQueue</b> property of the ConnectionFactory object is used instead.
brokerPubQueueManager <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• "" (<b>empty string</b>)</li> <li>• A queue manager name</li> </ul>	The name of the queue manager that owns the queue where messages published on the topic are sent.

Table 164. Properties that are specific to a Topic object (continued)

Name of property	Type	Valid values (default value in bold)	Description
brokerVersion <sup>1</sup>	String	<ul style="list-style-type: none"> <li>• <b>Not set</b></li> <li>• 1</li> <li>• 2</li> </ul>	The version of the broker being used. The value of this property overrides the value of the <b>brokerVersion</b> property of the ConnectionFactory object. However, if you do not set the value of this property, the value of the <b>brokerVersion</b> property of the ConnectionFactory object is used instead.

**Note:**

1. This property can be used with Version 7.0 of IBM WebSphere MQ classes for JMS but does not affect an application connected to a Version 7.0 queue manager unless the providerVersion property of the ConnectionFactory object is set to a version number less than 7.

The following example shows a set of properties of a Queue object:

```
expiry: UNLIM
persistence: QDEF
baseQueueManagerName: ExampleQM
baseQueueName: SYSTEM.JMS.TEMPQ.MODEL
```

The following example shows a set of properties of a Topic object:

```
expiry: UNLIM
persistence: NON
baseTopicName: myTestTopic
```

**Related information:**

- Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client
- Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows
- Configuring JMS resources in WebSphere Application Server

**Verifying the resource adapter installation**

The installation verification test (IVT) program for the IBM MQ resource adapter is supplied as an EAR file. To use the program, you must deploy it and define some objects as JCA resources.

**About this task**

The installation verification test (IVT) program is supplied as an enterprise archive (EAR) file called `wmq.jmsra.ivt.ear`. This file is installed with IBM MQ classes for JMS in the same directory as the IBM MQ resource adapter RAR file, `wmq.jmsra.rar`. For information about where these files are installed, see “Installing the IBM MQ resource adapter” on page 1099.

You must deploy the IVT program on your application server. The IVT program includes a servlet and an MDB that tests that a message can be sent to, and received from, an IBM MQ queue. You can use the IVT program to verify that the IBM MQ resource adapter has been correctly configured to support distributed transactions. If you are deploying the IBM MQ resource adapter in a non-IBM application server, IBM Service might ask you to demonstrate the IVT working to validate that your application server is correctly configured.

Before you can run the IVT program, you must define a ConnectionFactory object, a Queue object and possibly an Activation Specification object as JCA resources, and ensure that your application server

creates JMS objects from these definitions and binds them into a JNDI namespace. You can choose the properties of the objects to match the host and port settings of your own QueueManager, but the following set of properties is a simple example taken from WebSphere Application Server Liberty:

```
<!-- IVT Connection factory -->
<jmsQueueConnectionFactory connectionManagerRef="ConMgrIVT" jndiName="IVTCF">
 <properties.wmqJms channel="SYSTEM.DEF.SVRCONN" hostname="localhost" port="1550" transportType="CLIENT"/>
</jmsQueueConnectionFactory>
<connectionManager id="ConMgrIVT" maxPoolSize="10"/>

<!-- IVT Queues -->
<jmsQueue id="IVTQueue" jndiName="IVTQueue">
 <properties.wmqJms baseQueueName="RA_IVT_Q"/>
</jmsQueue>
```

By default, the IVT program expects a ConnectionFactory object to be bound in the JNDI namespace with the name `jms/ivt/IVTCF` and a Queue object to be bound with the name `jms/ivt/IVTQueue`. You can use different names, but if you do, you must enter the names of the objects on the initial page of the IVT program and modify the EAR file appropriately.

After you have deployed the IVT program, and the application server has created the JMS objects and bound them into the JNDI namespace, you can start the IVT program by completing the following steps.

## Procedure

1. Start the IVT program by entering a URL in the following format into your web browser:

```
http://app_server_host:port/wmq_ivt/
```

where `app_server_host` is the IP address or host name of the system on which your application server is running, and `port` is the number of the TCP port on which the application server is listening. Here is an example:

```
http://localhost:9080/wmq_ivt/
```

Figure 146 shows the initial page of the IVT program.



**IBM MQ JavaEE 7 Connector Architecture IVT**

**Installation Verification Test**  
Check to ensure that the IBM MQ J2EE Connector Architecture resource adapter is correctly installed.

Connection Factory:

Destination:

Figure 146. The initial page of the IVT program

2. To run the test, click **Run IVT**.

Figure 147 on page 1137 shows the page that is displayed if the IVT is successful.



# IBM MQ JavaEE 7 Connector Architecture IVT

## Running Installation Verification Test:

Using Connection Factory:*IVTCF*  
Using Destination:*IVTQueue*

Creating initial context...	☑
Looking up MQ Connection Factory...	☑
Looking up Destination...	☑
Creating connection...	☑
Starting connection...	☑
Creating session...	☑
Creating a temporary reply queue...	☑
Creating message consumer...	☑
Creating message producer...	☑
Creating message...	☑
Sending message to the MDB...	☑
Receiving response message from the MDB...	☑
Closing connection...	☑

## Installation Verification Test completed successfully!

[View Message Contents](#)

[Re-run Installation Verification Test](#)

Figure 147. Page showing the results of a successful IVT

If the IVT fails, a page like that shown in Figure 148 on page 1138 is displayed. To obtain further information about the cause of the failure, click **View Stack Trace**.

# IBM WebSphere MQ J2EE Connector Architecture IVT

## Running Installation Verification Test:

Using Connection Factory:java:comp/env/IVTCF  
Using Destination:java:comp/env/IVTQueue

Creating initial context...	☒
Looking up MQ Connection Factory...	☒
Looking up Destination...	☒
Creating connection...	☒
Starting connection...	☒
Creating session...	☒
Creating a temporary reply queue...	☒
Creating message consumer...	☒
Creating message producer...	☒
Creating message...	☒
Sending message to the MDB... failed to send message!	☒

## Installation Verification Test failed!

Error received - JMS Exception:

```
com.ibm.msg.client.jms.DetailedIllegalStateException: JMSWMQ2007: Failed to send a message to destination 'TEST.QUEUE'.
```

JMS attempted to perform an MQPUT or MQPUT1; however WebSphere MQ reported an error.

Use the linked exception to determine the cause of this error.

[View Stack Trace](#)

## Installation Verification Test failed!

[Retry Installation Verification Test](#)

[Change IVT parameters](#)

Figure 148. Page showing the results of an IVT that failed

## Installing and testing the resource adapter in GlassFish Server

### Windows

To install the IBM MQ resource adapter in GlassFish Server on a Windows operating system, you must first create and start a domain. You can then deploy and configure the resource adapter, and deploy and run the installation verification test (IVT) application.

### About this task

**Important:** These instructions are for GlassFish Server version 4.

This task assumes that you have a running GlassFish Server application server, and that you are familiar with standard administration tasks for it. This task also assumes that you have an IBM MQ installation on your local system and that you are familiar with standard administration tasks.

**Note:** In order to complete the following task steps, you must have a functioning IBM MQ installation, with the following objects configured:

- A queue manager called QM, that is started on port 1414, that uses channel SYSTEM.DEF.SVRCONN, and that connects using Client transport.
- A queue called Q1.

## Procedure

1. Start the GlassFish Server **asadmin** shell program.
  - a. Open the Windows command line and navigate to the *GlassFish/bin* directory, where *GlassFish* is the directory where GlassFish Server version 4 is installed.
  - b. Enter the command **asadmin** in the command line. The **asadmin** command opens a shell program in the command line that enables you to create a new domain.

GlassFish Server version 4 is started on your system.

2. Create, and then start a domain.
  - a. Use the **create-domain** command, specifying the port and domain name, to create a new domain. Enter the following command on the command line:  
`create-domain --adminport port domain_name`

where *port* is the port number, and *domain\_name* is the name you want the domain to use.

**Note:** The **create-domain** command has many optional parameters associated with it. However, for this task you need the `--adminport` parameter only. For more information, see the product documentation for GlassFish Server version 4.

If the port that you specified is in use, the following message appears:

```
Port for domain_name port is in use
```

If the domain name you specified is in use, you receive a message telling you that your specified name is already in use, as well as a list of all domain names that are currently unavailable.

- b. When prompted to input a user name and password, enter the credentials to be used to log on to the application server through a web browser. If the command completes successfully, a message summarizing the domain creation is displayed on the command line, including the message `Command create-domain executed successfully.`  
You have successfully created a domain.
- c. Start your domain by entering the following command into the command line:  
`start-domain domain_name`

where *domain\_name* is the domain name you previously specified.

3. Use a web browser to access GlassFish application server.
  - a. In the address bar of a web browser enter the following command:  
`localhost:port`  
where *port* is the port that you specified earlier when creating your domain. The GlassFish Console is displayed.
  - b. When the GlassFish Console has loaded, and you are prompted for a user name and password, enter the credentials that you specified in step 2b.
4. Upload the resource adapter to GlassFish Server 4.
  - a. On the toolbar **Common Tasks** select the **Applications** menu item to display the Applications page.
  - b. Click the **Deploy** button, to open the Deploy Applications or Modules page.
  - c. Click the **Browse** button to and then navigate to the location of the `wmq.jmsra.rar` file. Select the file and click **OK**.
5. Create a connection pool.
  - a. On the toolbar, under **Resources**, select the **Connectors** menu item.
  - b. Then select **Connector Connection Pools** menu item, to open the **Connector Connection Pools** page.
  - c. Click **New** to open the New Connector Connection Pool (Step 1 of 2) page.

- d. On New Connector Connection Pool (Step 1 of 2) page, input the pool name as `jms/ivt/IVTCF-Connection-Pool` into the **Pool Name** field.
  - e. In the **Resource Adapter** field select **wmq.jmsra**.
  - f. In the **Connection Definition** field enter `javax.jms.ConnectionFactory`.
  - g. Select **Next**, then select **Finish**.
6. Create the connector resources.
- a. On the toolbar, under the **Connectors** menu, select the **Connector Resource** option, to open the Connector Resources page.
  - b. Select **New**, to open the **New Connector Resource** page.
  - c. In the **JNDI Name** field, enter `IVTCF`.
  - d. In the **Pool Name** field, enter `jms/ivt/IVTCF-Connection-Pool`.
  - e. Leave all other fields empty.
  - f. For each of the following property/value pairs, click **Add Property**, and enter the property name and the value as shown in the following example:
    - name: host; value: localhost
    - name: port; value 1414
    - name: channel; value: SYSTEM.DEF.SVRCONN
    - name: queueManager; value: QM
    - name: transportType; value: CLIENT

**Note:** Make sure that you use the correct values for your own configuration settings, which might be different from the ones shown in this example.
  - g. In the toolbar, under **Connectors**, select the **Admin Object Resources** menu item, to open the **Admin Object Resources** page.
  - h. In the **Admin Object Resources** page, click **New** to open the **New Admin Object Resource** page.
  - i. In the **JNDI Name** field, enter `IVTQueue`.
  - j. In the **Resource Adapter** field, enter `wmq.jmsra`.
  - k. In the **Resource Type** field, enter `javax.jms.Queue`.
  - l. Leave the **Class Name** field as it is.
  - m. For each of the following property/value pairs, click **Add Property**, and enter the property name and the value as shown in the following example:
    - name: name; value: IVTQueue
    - name: baseQueueManagerName; value QM
    - name: baseQueueName; value: Q1

**Note:** Make sure that you use the correct values for your own configuration settings, which might be different from the ones shown in this example.
  - n. Click **OK**.
  - o. Select the **Enabled** checkbox, then click **Enable**.
7. Deploy the EAR file `wmq.jmsra.ivt.ear` into GlassFish Server.
- a. Click the **Applications** option in the toolbar to display the Applications page.
  - b. Click **Deploy** to add the IVT application.
  - c. In the **Location** field navigate to, and select, the `wmq.jmsra.ivt.ear`.
  - d. In the **Virtual Servers** field, select **server**, and then click **OK**.
8. Launch the IVT program.
- a. Click the **Applications** option in the toolbar to display the Applications page.
  - b. Click on `wmq.jmsra.ivt` in the Deployed Applications table.

- c. Click the **Launch** button, in the Modules and Components table.
- d. Select the http: link.
- e. Click **Run IVT**.

You have launched the IVT program, and if you are successful, the following output is displayed:

## Running Installation Verification Test:

Using Connection Factory:*IVTCF*  
Using Destination:*IVTQueue*

Creating initial context...	✓
Looking up MQ Connection Factory...	✓
Looking up Destination...	✓
Creating connection...	✓
Starting connection...	✓
Creating session...	✓
Creating a temporary reply queue...	✓
Creating message consumer...	✓
Creating message producer...	✓
Creating message...	✓
Sending message to the MDB...	✓
Receiving response message from the MDB...	✓
Closing connection...	✓

## Installation Verification Test completed successfully!

[View Message Contents](#)

[Re-run Installation Verification Test](#)

Figure 149. Successful IVT output

## Installing and testing the resource adapter in Wildfly

V 9.0.2

V 9.0.0.1

If you are installing the IBM MQ resource adapter in Wildfly V10, you must first make some configuration file changes to add a subsystem definition for the IBM MQ resource adapter. You can then deploy the resource adapter and test it by installing and running the installation verification test (IVT) application.

### About this task

**Important:** These instructions are for Wildfly V10.

This task assumes that you have a running WildFly application server, and that you are familiar with standard administration tasks for it. This task also assumes that you have an IBM MQ installation and that you are familiar with standard administration tasks.

### Procedure

1. Create an IBM MQ queue manager called ExampleQM, and set it up as described in “Configuring a queue manager to accept client connections on Windows, UNIX and IBM i” on page 470. When setting up the queue manager, note the following points:

- The listener must be started on port 1414.
- The channel to be used is called SYSTEM.DEF.SVRCONN.
- The queue used by the IVT application is named TEST.QUEUE.

The model queue SYSTEM.DEFAULT.MODEL.QUEUE also needs to be granted DSP and PUT authority so this application can create a temporary reply queue.

2. Edit the configuration file *WildFly\_Home/standalone/configuration/standalone-full.xml* and add the following subsystem:

```
<subsystem xmlns="urn:jboss:domain:resource-adapters:4.0">
 <resource-adapters>
 <resource-adapter id="wmq.jmsra">
 <archive>
 wmq.jmsra.rar
 </archive>
 <transaction-support>NoTransaction</transaction-support>
 <connection-definitions>
 <connection-definition class-name="com.ibm.mq.connector.outbound.ManagedConnectionFactoryImpl"
 jndi-name="java:jboss/jms/ivt/IVTCF" enabled="true" use-java-context="true"
 pool-name="IVTCF">
 <config-property name="channel">SYSTEM.DEF.SVRCONN
 </config-property>
 <config-property
 name="hostName">localhost
 </config-property>
 <config-property name="transportType">
 CLIENT
 </config-property>
 <config-property name="queueManager">
 ExampleQM
 </config-property>
 <config-property name="port">
 1414
 </config-property>
 </connection-definition>
 <connection-definition class-name="com.ibm.mq.connector.outbound.ManagedConnectionFactoryImpl"
 jndi-name="java:jboss/jms/ivt/JMS2CF" enabled="true" use-java-context="true"
 pool-name="JMS2CF">
 <config-property name="channel">
 SYSTEM.DEF.SVRCONN
 </config-property>
 </connection-definition>
 </connection-definitions>
 </resource-adapter>
 </resource-adapters>
</subsystem>
```

```

 <config-property name="hostName">
 localhost
 </config-property>
 <config-property name="transportType">
 CLIENT
 </config-property>
 <config-property name="queueManager">
 ExampleQM
 </config-property>
 <config-property name="port">
 1414
 </config-property>
 </connection-definition>
</connection-definitions>
<admin-objects>
 <admin-object class-name="com.ibm.mq.connector.outbound.MQQueueProxy"
 jndi-name="java:jboss/jms/ivt/IVTQueue" pool-name="IVTQueue">
 <config-property name="baseQueueName">
 TEST.QUEUE
 </config-property>
 </admin-object>
</admin-objects>
</resource-adapter>
</resource-adapters>
</subsystem>

```

3. Deploy the resource adapter to your server by copying the `wmq.jmsra.rar` file into the directory `WildFly_Home/standalone/deployments`.
4. Deploy the IVT application by copying the `wmq.jmsra.ivt.ear` file into the directory `WildFly_Home/standalone/deployments`.
5. Start the application server, by bringing up a command prompt, navigating to the directory `WildFly_Home/bin` and running the command:  
`standalone.bat -c standalone-full.xml`
6. Run the IVT application. For more information, see “Verifying the resource adapter installation” on page 1135. For Wildfly, the default URL is `http://localhost:8080/WMQ_IVT/`.

## Using the IBM MQ Headers package

The IBM MQ Headers package provides a set of helper interfaces and classes that you can use to manipulate the IBM MQ headers of a message. Typically, you use the IBM MQ Headers package because you want to perform administrative services by using the command server (by using Programmable Command Format (PCF) messages).

### About this task

The IBM MQ Headers package is located in the `com.ibm.mq.headers` and `com.ibm.mq.pcf` packages. You can use this facility for both of the two alternative APIs that IBM MQ provides for use in Java applications:

- IBM MQ classes for Java (also referred to as IBM MQ Base Java).
- IBM MQ classes for Java Message Service (IBM MQ classes for JMS, also referred to as IBM MQ JMS).

IBM MQ Base Java applications typically manipulate `MQMessage` objects, and the Headers support classes can directly interact with these objects, since they natively understand the IBM MQ Base Java interfaces.

In IBM MQ JMS, the payload for a message is typically a `String` or a byte array object, which can be manipulated with `DataInput` and `DataOutput` streams. The IBM MQ Headers package can be used to interact with these data streams and is suitable for manipulating any MQ messages that are sent and received by IBM MQ JMS applications.

Therefore, although the IBM MQ Headers package contains references to the IBM MQ Base Java package, it is also intended for use within IBM MQ JMS applications and is suitable for use within Java Platform, Enterprise Edition (Java EE) environments.

A typical way in which you might use the IBM MQ Headers package is to manipulate administration messages in Programmable Command Format (PCF), for example for any of the following reasons:

- To access details about an IBM MQ resource.
- To monitor the depth of a queue.
- To inhibit access to a queue.

By using PCF messages with the IBM MQ JMS API, this kind of administration of application-centric resources can be performed from within Java EE applications without having to resort to using the IBM MQ Base Java API.

## Procedure

- To use the IBM MQ Headers package to manipulate message headers for IBM MQ classes for Java, see “Using with IBM MQ classes for Java.”
- To use the IBM MQ Headers package to manipulate message headers for IBM MQ classes for JMS, see “Using with IBM MQ classes for JMS” on page 1145.

## Using with IBM MQ classes for Java

IBM MQ classes for Java applications typically manipulate MQMessage objects, and the Headers support classes can directly interact with these objects, since they natively understand the IBM MQ classes for Java interfaces.

## About this task

IBM MQ provides some sample applications that demonstrate how to use the IBM MQ Headers package with the IBM MQ Base Java API (IBM MQ classes for Java).

The samples show two things:

- How to create a PCF message to perform an administrative action and parse the response message.
- How to send this PCF message using the IBM MQ classes for Java.

Depending on which platform you are using, these samples are installed under the `pcf` directory in the `samples` or `tools` directory of your IBM MQ installation (see “Installation directories for IBM MQ classes for Java” on page 1028).

## Procedure

1. Create a PCF message to perform an administrative action and parse the response message.
2. Send this PCF message using the IBM MQ classes for Java.



## Related concepts:

“Handling IBM MQ message headers with IBM MQ classes for Java” on page 1053

Java classes are provided representing different types of message header. Two helper classes are also provided.

“Handling PCF messages with IBM MQ classes for Java” on page 1059

Java classes are provided to create and parse PCF-structured messages, and to facilitate sending PCF requests and collecting PCF responses.

## Using with IBM MQ classes for JMS

To use the IBM MQ Headers with the IBM MQ classes for JMS, you carry out the same essential steps as for IBM MQ classes for Java. The PCF message can be created and the response parsed in exactly the same way by using the IBM MQ Headers package and the same sample code as for IBM MQ classes for Java.

## About this task

To send a PCF message using the IBM MQ API, the message payload must be written into a JMS Bytes Message, and sent using the standard JMS APIs. The only consideration is that the message must not contain a JMS RFH2 or any other headers with specific values in the MQMD.

To send a PCF message, complete the following steps. The way in which the PCF message is created, and information is extracted from the response message is the same as for IBM MQ classes for Java (see “Using with IBM MQ classes for Java” on page 1144).

## Procedure

1. Create a JMS Queue Destination that represents the SYSTEM.ADMIN.COMMAND.QUEUE. IBM MQ JMS applications send the PCF messages to the SYSTEM.ADMIN.COMMAND.QUEUE, and need access to a JMS Destination object that represents this queue. The Destination must have the following properties set:

```
WMQ_MQMD_WRITE_ENABLED = YES
WMQ_MESSAGE_BODY = MQ
```

If you are using WebSphere Application Server, you must define these properties as custom properties on the Destination.

To create the destination programmatically from within an application, use the following code:

```
Queue q1 = session.createQueue("SYSTEM.ADMIN.COMMAND.QUEUE");
((MQQueue) q1).setIntProperty(WMQConstants.WMQ_MESSAGE_BODY,
 WMQConstants.WMQ_MESSAGE_BODY_MQ);
((MQQueue) q1).setMQMDWriteEnabled(true);
```

2. Convert a PCF message into a JMS Bytes message containing the correct MQMD values. A JMS Bytes message needs to be created, and the PCF Message written to it. A response queue needs to be created, but this needs to have no specific settings.

The following sample code snippet shows how to create a JMS Bytes Message and write a com.ibm.mq.headers.pcf.PCFMessage object into it. The PCFMessage object (pcfCmd) has previously been built using the IBM MQ Headers package. (Note the package to load the PCFMessage is com.ibm.mq.headers.pcf.PCFMessage).

```
// create the JMS Bytes Message
final BytesMessage msg = session.createBytesMessage();

// Create the wrapping streams to put the bytes into the message payload
ByteArrayOutputStream baos = new ByteArrayOutputStream();
DataOutput dataOutput = new DataOutputStream(baos);

// Set the JMSReplyTo so the answer comes back
msg.setJMSReplyTo(new MQQueue("adminResp"));
```

```

// write the pcf into the stream
pcfCmd.write(dataOutput);
baos.flush();
msg.writeBytes(baos.toByteArray());

// we have taken control of the MD, so need to set all
// flags in the MD that we require - main one is the format
msg.setJMSPriority(4);
msg.setIntProperty(WMQConstants.JMS_IBM_MQMD_PERSISTENCE,
 CMQC.MQPER_NOT_PERSISTENT);
msg.setIntProperty(WMQConstants.JMS_IBM_MQMD_EXPIRY, 300);
msg.setIntProperty(WMQConstants.JMS_IBM_MQMD_REPORT,
 CMQC.MQRO_PASS_CORREL_ID);
msg.setStringProperty(WMQConstants.JMS_IBM_MQMD_FORMAT, "MQADMIN");

// and send the message
sender.send(msg);

```

3. Send the message, and receive the response using the standard JMS APIs.
4. Convert the response message into a PCF message for processing. To retrieve the response message and process it as a PCF message, use the following code:

```

// Get the message back
BytesMessage msg = (BytesMessage) consumer.receive();

// get the size of the bytes message & read into an array
int bodySize = (int) msg.getBodyLength();
byte[] data = new byte[bodySize];
msg.readBytes(data);

// Read into Stream and DataInput Stream
ByteArrayInputStream bais = new ByteArrayInputStream(data);
DataInput dataInput = new DataInputStream(bais);

// Pass to PCF Message to process
PCFMessage response = new PCFMessage(dataInput);

```

#### Related concepts:

“JMS messages” on page 875

JMS messages are composed of a header, properties, and a body. JMS defines five types of message body.

## Setting up IBM MQ on IBM i with Java and JMS



This collection of topics gives an overview of how you set up and test IBM MQ with Java and JMS on IBM i using CL commands or the qshell environment.

**Note:** From IBM MQ Version 8.0, ldap.jar, jndi.jar and jta.jar are part of the JDK.

### Using CL commands

The CLASSPATH that you set, is for testing with MQ base Java, JMS with JNDI, and JMS without JNDI.

If you do not use a .profile file under your /home/Userprofile directory, you will need to set the system level environment variables below. You can check to see if they are set using the **WRKENVVAR** command.

1. To view the environment variables for the entire system issue the command: **WRKENVVAR LEVEL(\*SYS)**
2. To view the environment variables specific to your job issue the command : **WRKENVVAR LEVEL(\*JOB)**
3. If the CLASSPATH is not set, do the following:

```
ADDENVVAR ENVVAR(CLASSPATH)
VALUE('.:QIBM/ProdData/mqm/java/lib/com.ibm.mq.jar
:QIBM/ProdData/mqm/java/lib/connector.jar:QIBM/ProdData/mqm/java/lib
:QIBM/ProdData/mqm/java/samples/base
:QIBM/ProdData/mqm/java/lib/com.ibm.mqjms.jar
:QIBM/ProdData/mqm/java/lib/jms.jar
:QIBM/ProdData/mqm/java/lib/providerutil.jar
:QIBM/ProdData/mqm/java/lib/fscontext.jar:') LEVEL(*SYS)
```

4. If QIBM\_MULTI\_THREADED is not set, issue the following command:

```
ADDENVVAR ENVVAR(QIBM_MULTI_THREADED) VALUE('Y') LEVEL(*SYS)
```

5. If QIBM\_USE\_DESCRIPTOR\_STDIO is not set, issue the following command:

```
ADDENVVAR ENVVAR(QIBM_USE_DESCRIPTOR_STDIO) VALUE('I') LEVEL(*SYS)
```

6. If QSH\_REDIRECTION\_TEXTDATA is not set, issue the following command:

```
ADDENVVAR ENVVAR(QSH_REDIRECTION_TEXTDATA) VALUE('Y') LEVEL(*SYS)
```

## Using the qshell environment

If you use the QSHELL environment, you can set up a .profile in your /home/Userprofile directory. For more information reference the Qshell Interpreter (qsh) documentation.

Specify the following in the .profile. Note that the CLASSPATH statement must be on a single line, or separated onto different lines using the \ character as shown.

```
CLASSPATH=.:QIBM/ProdData/mqm/java/lib/com.ibm.mq.jar: \
/QIBM/ProdData/mqm/java/lib/connector.jar: \
/QIBM/ProdData/mqm/java/lib: \
/QIBM/ProdData/mqm/java/samples/base: \
/QIBM/ProdData/mqm/java/lib/com.ibm.mqjms.jar: \
/QIBM/ProdData/mqm/java/lib/jms.jar: \
/QIBM/ProdData/mqm/java/lib/providerutil.jar: \
/QIBM/ProdData/mqm/java/lib/fscontext.jar:
HOME=/home/XXXXX
LOGNAME=XXXXX
PATH=/usr/bin:
QIBM_MULTI_THREADED=Y QIBM_USE_DESCRIPTOR_STDIO=I
QSH_REDIRECTION_TEXTDATA=Y
TERMINAL_TYPE=5250
```

Ensure that the QMQMJAVA library is in the library list by issuing the command **DSPLIBL**.

If the QMQMJAVA library is not in the list, add it using the following command: **ADDLIBLE LIB(QMQMJAVA)**

## Testing IBM MQ on IBM i with Java



How you test IBM MQ with Java using the MQIVP sample program.

### Testing IBM MQ base Java

Carry out the following procedure:

1. Verify that the queue manager is started, and that the state of the queue manager is ACTIVE, by issuing the following command:

```
WRKMQM MQMNAME(QMGRNAME)
```

2. Verify that the JAVA.CHANNEL server connection channel has been created by issuing the following command:

```
WRKMQMCHL CHLNAME(JAVA.CHANNEL) CHLTYPE(*SVRCN) MQMNAME(QMGRNAME)
```

- a. If the JAVA.CHANNEL does not exist, issue the following command:

```
CRTMQMCHL CHLNAME(JAVA.CHANNEL) CHLTYPE(*SVRCN) MQMNAME(QMGRNAME)
```

3. Verify the queue manager listener is running for port 1414 or whichever port you are using, by issuing the **WRKMQMLSR** command.

- a. If no listener has been started for the queue manager, issue the following command:

```
STRMQMLSR PORT(XXXX) MQMNAME(QMGRNAME)
```

### Running the MQIVP sample test program

1. Start the qshell, from the command line by issuing the command STRQSH
2. Verify that the correct CLASSPATH is set by issuing the **export** command, and then issue the **cd** command as follows:

```
cd /qibm/proddata/mqm/java/samples/wmqjava/samples
```

3. Run the **java** program by issuing the following command:

```
java MQIVP
```

You can press the ENTER key when prompted for:

- Type of connection
- IP address
- Queue manager name

to use the default values. This verifies the product bindings, which can be found in the QMQMJAVA library.

You receive output similar to the following example. Note that the copyright statement depends upon the version of the product that you are using.

```
> java MQIVP
MQSeries for Java Installation Verification Program
5724-H72 (C) Copyright IBM Corp. 2011. All Rights Reserved.
=====
```

```
Please enter the IP address of the MQ server :
```

```
>
```

```
Please enter the queue manager name :
```

```
>
```

```
Attaching Java program to QIBM/ProdData/mqm/java/lib/connector.JAR.
```

```
Success: Connected to queue manager.
```

```
Success: Opened SYSTEM.DEFAULT.LOCAL.QUEUE
```

```
Success: Put a message to SYSTEM.DEFAULT.LOCAL.QUEUE
```

```
Success: Got a message from SYSTEM.DEFAULT.LOCAL.QUEUE
```

```
Success: Closed SYSTEM.DEFAULT.LOCAL.QUEUE
```

Success: Disconnected from queue manager

```
Tests complete -
SUCCESS: This MQ Transport is functioning correctly.
Press Enter to continue ...
>
$
```

## Testing IBM MQ Java client connection

You must specify the:

- Connection type
- IP address
- Port
- Server connection channel
- Queue manager

You receive output similar to the following example. Note that the copyright statement depends upon the version of the product that you are using.

```
> java MQIVP
MQSeries for Java Installation Verification Program
5724-H72 (C) Copyright IBM Corp. 2011. All Rights Reserved.
=====
```

```
Please enter the IP address of the MQ server :
> x.xx.xx.xx
Please enter the port to connect to : (1414)
> 1470
Please enter the server connection channel name :
> JAVA.CHANNEL
Please enter the queue manager name :
> KAREN01
Success: Connected to queue manager.
Success: Opened SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Put a message to SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Got a message from SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Closed SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Disconnected from queue manager
```

```
Tests complete -
SUCCESS: This MQ Transport is functioning correctly.
Press Enter to continue ...
>
$
```

## Testing IBM MQ on IBM i with JMS



How you test IBM MQ with JMS with and without JNDI

### Testing JMS without JNDI using the IVTRun sample

Carry out the following procedure:

1. Verify that the queue manager is started, and that the state of the queue manager is ACTIVE, by issuing the following command:

```
WRKMQM MQMNAME(QMGRNAME)
```

2. Start the qshell, from the command line, by issuing the **STRQSH** command.

3. Use the **cd** command to change directory as follows:

```
cd /qibm/proddata/mqm/java/bin
```

4. Run the script file:

```
IVTRun -nojndi [-m qmgrname]
```

You receive output similar to the following example. Note that the copyright statements depend upon the versions of the products that you are using:

```
> IVTRun -nojndi -m ELCRTP19
```

```
Attaching Java program to
/QIBM/ProdData/mqm/java/lib/com.ibm.mqjms.JAR.
```

```
Attaching Java program to
/QIBM/ProdData/mqm/java/lib/jms.JAR.
```

```
5724-H72, 5724-B41, 5655-F10 (c) Copyright IBM Corp. 2011.
All Rights Reserved.
WebSphere MQ classes for Java(tm) Message Service 5.300
Installation Verification Test
```

```
Creating a QueueConnectionFactory
Creating a Connection
Creating a Session
Creating a Queue
Creating a QueueSender
Creating a QueueReceiver
Creating a TextMessage
Sending the message to SYSTEM.DEFAULT.LOCAL.QUEUE
Reading the message back again
```

```
Got message:
JMS Message class: jms_text
JMSType: null
JMSDeliveryMode: 2
JMSExpiration: 0
JMSPriority: 4
JMSMessageID: ID:c1d4d840c5d3c3d9e3d7f1f9404040403ccf041f0000c012
JMSTimestamp: 1020273404500
JMSCorrelationID:null
JMSDestination: queue:///SYSTEM.DEFAULT.LOCAL.QUEUE
JMSReplyTo: null
JMSRedelivered: false
JMS_IBM_PutDate:20040326
JMSXAppID:QP0ZSPWT STANLEY 170302
JMS_IBM_Format:MQSTR
JMS_IBM_PutAppIType:8
JMS_IBM_MsgType:8
JMSXUserID:STANLEY
JMS_IBM_PutTime:13441354
```

```

JMSXDeliveryCount:1
A simple text message from the MQJMSIVT program
Reply string equals original string
Closing QueueReceiver
Closing QueueSender
Closing Session
Closing Connection
IVT completed OK
IVT finished
$
>
$

```

## Testing IBM MQ JMS client mode without JNDI

Carry out the following procedure:

1. Verify that the queue manager is started, and that the state of the queue manager is ACTIVE, by issuing the following command:  
**WRKMQM MQMNAME(QMGRNAME)**
2. Verify that the server connection channel is created, by issuing the following command:  
**WRKMQMCHL CHLNAME( SYSTEM.DEF.SVRCONN ) CHLTYPE(\*SVRCN)  
MQMNAME(QMGRNAME)**
3. Verify that the listener is started for the correct port, by issuing the **WRKMQLSR** command
4. Start the qshell, from the command line, by issuing the **STRQSH** command.
5. Verify that the CLASSPATH is correct by issuing the **export** command.
6. Use the **cd** command to change directory as follows:  
**cd /qibm/proddata/mqm/java/bin**
7. Run the script file:  
**IVTRun -nojndi -client -m QMgrName -host hostname [-port port] [-channel channel]**

You receive output similar to the following example. Note that the copyright statements depend upon the versions of the products that you are using.

```
> IVTRun -nojndi -client -m ELCRTP19 -host ELCRTP19 -port 1414 -channel SYSTEM.DEF.SVRCONN
```

```

5724-H72, 5724-B41, 5655-F10 (c) Copyright IBM Corp. 2011.
All Rights Reserved.
WebSphere MQ classes for Java(tm) Message Service 5.300
Installation Verification Test

```

```

Creating a QueueConnectionFactory
Creating a Connection
Creating a Session
Creating a Queue
Creating a QueueSender
Creating a QueueReceiver
Creating a TextMessage
Sending the message to SYSTEM.DEFAULT.LOCAL.QUEUE
Reading the message back again

```

```

Got message:
JMS Message class: jms_text
JMSType: null
JMSDeliveryMode: 2
JMSExpiration: 0
JMSPriority: 4
JMSMessageID: ID:c1d4d840c5d3c3d9e3d7f1f9404040403ccf041f0000d012
JMSTimestamp: 1020274009970
JMSCorrelationID:null
JMSDestination: queue:///SYSTEM.DEFAULT.LOCAL.QUEUE

```

```

JMSReplyTo: null
JMSRedelivered: false
JMS_IBM_PutDate:20040326
JMSXAppID:MQSeries Client for Java
JMS_IBM_Format:MQSTR
JMS_IBM_PutAppIType:28
JMS_IBM_MsgType:8
JMSXUserID:QMQM
JMS_IBM_PutTime:14085237
JMSXDeliveryCount:1
A simple text message from the MQJMSIVT program
Reply string equals original string
Closing QueueReceiver
Closing QueueSender
Closing Session
Closing Connection
IVT completed OK
IVT finished
$

```

## Testing IBM MQ JMS with JNDI

Verify that the queue manager is started, and that the state of the queue manager is ACTIVE, by issuing the following command:

```
WRKMQM MQMNAME(QMGRNAME)
```

### Using the IVTRun sample test script

Carry out the following procedure:

1. Make the appropriate changes to the JMSAdmin.config file. To edit this file use the **EDTF** (Edit File) command from an IBM i command line
 

```
EDTF '/qibm/proddata/mqm/java/bin/JMSAdmin.config'
```

  - a. To use LDAP for Weblogic, remove the comment from:
 

```
INITIAL_CONTEXT_FACTORY=com.sun.jndi.ldap.LdapCtxFactory
```
  - b. To use LDAP for WebSphere Application Server, remove the comment from:
 

```
INITIAL_CONTEXT_FACTORY=com.ibm.ejs.ns.jndi.CNInitialContextFactory
```
  - c. To test the file system, remove the comment from:
 

```
INITIAL_CONTEXT_FACTORY=com.sun.jndi.fscontext.RefFSContextFactory
```
  - d. Ensure that you have selected the correct PROVIDER\_URL, by removing the comment from the appropriate line.
  - e. Comment out all other lines using the # symbol.
  - f. When you have completed all your changes, press **F2=Save** and **F3=Exit**.
2. Start the qshell, from the command line, by issuing the **STRQSH** command.
3. Verify that the CLASSPATH is correct by issuing the **export** command.
4. Use the **cd** command to change directory as follows:
 

```
cd /qibm/proddata/mqm/java/bin
```
5. Start the **IVTSetup** script to create the administered objects (*MQQueueConnectionFactory* and *MQQueue*), by issuing the **IVTSetup** command.
6. Run the IVTRun script by issuing the following command:
 

```
IVTRun -url providerURL [-icf initCtxFact]
```

You receive output similar to the following example. Note that the copyright statements depend upon the versions of the products that you are using.



```

> IVTSetup
+ Creating script for object creation within JMSAdmin
+ Calling JMSAdmin in batch mode to create objects
Ignoring unknown flag: -i

5724-H72 (c) Copyright IBM Corp. 2011. All Rights Reserved.
Starting WebSphere MQ classes for Java(tm) Message Service Administration

InitCtx>
InitCtx>
InitCtx>
InitCtx>
InitCtx>
Stopping MQSeries classes for Java(tm) Message Service Administration

+ Administration done; tidying up files
+ Done!
$

> IVTRun -url file:///tmp/mqjms -icf com.sun.jndi.fscontext.ReffSContextFactory

5724-H72 (c) Copyright IBM Corp. 2011. All Rights Reserved.
MQSeries classes for Java(tm) Message Service
Installation Verification Test

Using administered objects, please ensure that these are available

Retrieving a QueueConnectionFactory from JNDI
Creating a Connection
Creating a Session
Retrieving a Queue from JNDI
Creating a QueueSender
Creating a QueueReceiver
Creating a TextMessage
Sending the message to SYSTEM.DEFAULT.LOCAL.QUEUE
Reading the message back again

Got message:
JMS Message class: jms_text
JMSType: null
JMSDeliveryMode: 2
JMSExpiration: 0
JMSPriority: 4
JMSMessageID: ID:c1d4d840c5d3c3d9e3d7f1f9404040403ccf041f0000e012
JMSTimestamp: 1020274903770
JMSCorrelationID:null
JMSDestination: queue:///SYSTEM.DEFAULT.LOCAL.QUEUE
JMSReplyTo: null
JMSRedelivered: false
JMS_IBM_Format:MQSTR
JMS_IBM_PutApp1Type:8
JMSXDeliveryCount:1
JMS_IBM_MsgType:8
JMSXUserID:STANLEY
JMSXAppID:QP0ZSPWT STANLEY 170308
A simple text message from the MQJMSIVT program
Reply string equals original string
Closing QueueReceiver
Closing QueueSender
Closing Session
Closing Connection
IVT completed OK
IVT finished
$

```

---

## Developing applications for MQ Telemetry

Telemetry applications integrate sense and control devices with other sources of information available on the internet and in enterprises.

Develop applications for MQ Telemetry using design patterns, worked examples, sample programs, programming concepts, and reference information.

### Related information:

MQ Telemetry

Telemetry use cases

Installing MQ Telemetry

Administering MQ Telemetry

MQ Telemetry Reference

MQ Telemetry troubleshooting

## IBM MQ Telemetry Transport sample programs

Sample scripts are provided that work with a sample IBM MQ Telemetry Transport (MQTT) client application. Download the client application, then use the scripts to subscribe to a topic and publish a message.

### About this task

For further information, see the:

- Eclipse Paho web site, for the latest telemetry clients and an Android sample.
- MA9C: IBM Messaging Telemetry Clients SupportPac, which includes a similar set of clients to Paho, plus samples for Java and C.
- Deprecated MA9B: Mobile Messaging and M2M Client Pack includes a sample for iOS that was written in the IBM WebSphere MQ Version 7.5 time frame.

From these sites you can develop sample programs for publishing and subscribing IBM MQ Telemetry Transport and add security features.


### MQTTV3Sample program

Reference information about sample syntax and parameters for the MQTTV3Sample program.

#### Purpose

The MQTTV3Sample program can be used to publish a message and subscribe to a topic.

#### MQTT V3Sample syntax



#### Parameters

**-h** Print this help text and quit

- q Set quiet mode, instead of using the default mode of false.
- a Set publish or subscribe, instead of assuming the default action of publishing.
- t Publish or subscribe to topic, instead of publishing or subscribing to the default topic
- m Publish message text instead of sending the default publication text, "Hello from an MQTT v3 application".
- s Set QoS instead of using the default QoS, 2.
- b Connect to this host name or IP address instead of connecting to the default host name, localhost.
- p Use this port instead of using the default, 1883.

## Run the MQTTV3Sample program

To subscribe to a topic on Windows, use the command:

```
runMQTTV3Sample -a subscribe
```

To publish a message on Windows, use the command:

```
runMQTTV3Sample
```

For further information about running the sample scripts provided, see “IBM MQ Telemetry Transport sample programs” on page 1154.

## MQTT client programming concepts

The concepts described in this section help you to understand the client libraries for the MQTT protocol. The concepts complement the API documentation accompanying the client libraries.

For further information, see the:

- Eclipse Paho web site, for the latest telemetry clients and an Android sample.
- MA9C: IBM Messaging Telemetry Clients SupportPac, which includes a similar set of clients to Paho, plus samples for Java and C.
- Deprecated MA9B: Mobile Messaging and M2M Client Pack includes a sample for iOS that was written in the IBM WebSphere MQ Version 7.5 time frame.

To develop and run an MQTT client you need to copy or install these packages on the client device. You do not need to install a separate client runtime.

The licensing conditions for clients are associated with the server that you are connecting the clients to.

The MQTT client libraries are reference implementations of the MQTT protocol. You can implement your own clients in different languages suitable for different device platforms. See IBM MQ Telemetry Transport format and protocol.

The API documentation makes no assumptions about which MQTT server the client is connected to. The behavior of the client might differ slightly when connected to different servers. The descriptions that follow describe the behavior of the client when connected to the IBM MQ telemetry service.

## Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

### Callbacks

See the Eclipse Paho Web site for the latest changes to `MqttCallback`.

The `MqttCallback` interface has three callback methods:

#### **connectionLost(java.lang.Throwable cause)**

`connectionLost` is called when a communications error leads to the connection dropping. It is also called if the server drops the connection as a result of an error on the server after the connection has been established. Server errors are logged to the queue manager error log. The server drops the connection to the client, and the client calls `MqttCallback.connectionLost`.

The only remote errors thrown as exceptions on the same thread as the client application are exceptions from `MqttClient.connect`. Errors detected by the server after the connection is established are reported back to the `MqttCallback.connectionLost` callback method as throwables.

Typical server errors that result in `connectionLost` are authorization errors. For example, the telemetry server tries to publish on a topic on behalf of a client that is not authorized to publish on the topic. Anything that results in a `MQCC_FAIL` condition code being returned to the telemetry server can result in the connection being dropped.

#### **deliveryComplete(MqttDeliveryToken token)**

`deliveryComplete` is called by the MQTT client to pass a delivery token back to the client application; see “Delivery tokens” on page 1160. Using the delivery token, the callback can access the published message with the method `token.getMessage`.

When the application callback returns control to the MQTT client after being called by the `deliveryComplete` method, delivery is completed. Until delivery is completed, messages with QoS 1 or 2 are retained by the persistence class.

The call to `deliveryComplete` is a point of synchronization between the application and the persistence class. The `deliveryComplete` method is never called twice for the same message.

When the application callback returns from `deliveryComplete` to the MQTT client, the client calls `MqttClientPersistence.remove` for messages with QoS 1 or 2. `MqttClientPersistence.remove` deletes the locally stored copy of the published message.

From a transaction processing perspective, the call to `deliveryComplete` is a single phase transaction that commits the delivery. If processing fails during the callback, on restart of the client `MqttClientPersistence.remove` is called again to delete the local copy of the published message. The callback is not called again. If you are using the callback to store a log of delivered messages, you cannot synchronize the log with the MQTT client. If you want to store a log reliably, then update the log in the `MqttClientPersistence` class.

The delivery token and message are referenced by the main application thread and the MQTT client. The MQTT client dereferences the `MqttMessage` object when delivery is completed, and the delivery token object when the client disconnects. The `MqttMessage` object can be garbage collected after delivery is completed if the client application dereferences it. The delivery token can be garbage collected after the session is disconnected.

You can get `MqttDeliveryToken` and `MqttMessage` attributes after a message has been published. If you attempt to set any `MqttMessage` attributes after the message has been published the result is undefined.

The MQTT client continues to process delivery acknowledgments if the client reconnects to the previous session with the same `ClientIdentifier`; see “Clean sessions” on page 1158. The

MQTT client application must set `MqttClient.CleanSession` to `false` for the previous session, and set it to `false` in the new session. The MQTT client creates new delivery tokens and message objects in the new session for pending deliveries. It recovers the objects using the `MqttClientPersistence` class. If the application client still has references to the old delivery tokens and messages, dereference them. The application callback is called in the new session for any deliveries initiated in the previous session and completed in this session.

The application callback is called after the application client connects, when a pending delivery is completed. Before the application client connects, it can retrieve pending deliveries using the `MqttClient.getPendingDeliveryTokens` method.

Notice that the client application originally created the message object that is published, and its payload byte array. The MQTT client references these objects. The message object returned by the delivery token in the method `token.getMessage` is not necessarily the same message object created by the client. If a new MQTT client instance re-creates the delivery token, the `MqttClientPersistence` class re-creates the `MqttMessage` object. For consistency `token.getMessage` returns `null` if `token.isCompleted` is `true`, regardless of whether the message object was created by the application client or the `MqttClientPersistence` class.

### **messageArrived(MqttTopic topic, MqttMessage message)**

`messageArrived` is called when a publication arrives for the client that matched a subscription topic. `topic` is the publication topic, not the subscription filter. The two can be different if the filter contains wildcards.

If the topic matches multiple subscriptions created by the client, the client receives multiple copies of the publication. If a client publishes to a topic that it also subscribes to, it receives a copy of its own publication.

If a message is sent with a QoS of 1 or 2, the message is stored by the `MqttClientPersistence` class before the MQTT client calls `messageArrived`. `messageArrived` behaves like `deliveryComplete`: it is only called once for a publication, and the local copy of the publication is removed by `MqttClientPersistence.remove` when `messageArrived` returns to the MQTT client. The MQTT client drops its references to the topic and message when `messageArrived` returns to the MQTT client. The topic and message objects are garbage collected, if the application client has not held onto a reference to the objects.

## **Callbacks, threading, and client application synchronization**

The MQTT client calls a callback method on a separate thread to the main application thread. The client application does not create a thread for the callback, it is created by the MQTT client.

The MQTT client synchronizes callback methods. Only one instance of the callback method runs at a time. The synchronization makes it easy to update an object that tallies which publications have been delivered. One instance of the `MqttCallback.deliveryComplete` runs at a time, and so it is safe to update the tally without further synchronization. It is also the case that only one publication arrives at a time. Your code in the `messageArrived` method can update an object without synchronizing it. If you are referring to the tally, or the object that is being updated, in another thread, synchronize the tally or object.

The delivery token provides a synchronization mechanism between the main application thread and delivery of a publication. The method `token.waitForCompletion` waits until delivery of a specific publication is completed, or until an optional timeout expires. You might use `token.waitForCompletion` in the following way to process one publication at a time.

To synchronize with the `MqttCallback.deliveryComplete` method. Only when `MqttCallback.deliveryComplete` returns to the MQTT Client does `token.waitForCompletion` resume. Using this mechanism you can synchronize running code in `MqttCallback.deliveryComplete` before code runs in the main application thread.

What if you wanted to publish without waiting for each publication to be delivered, but want confirmation when all the publications have been delivered? If you publish on a single thread, the last publication to be sent is also the last to be delivered.

## Synchronization of requests sent to the server

Table 165 describes the methods in the MQTT Java client that send a request to the server. Unless the application client sets an indefinite timeout, the client never waits indefinitely for the server. If the client hangs, it is either an application programming problem, or a defect in the MQTT client.

Table 165. Synchronization behavior of methods that result in requests to the server

Method	Synchronization	Timeout interval
MqttClient.Connect	Waits for a connection to be established with the server.	Defaults to 30 seconds, or as set by a parameter, then throws an exception.
MqttClient.Disconnect	Waits for the MQTT client to finish any work it must do, and for the TCP/IP session to disconnect.	
MqttClient.Subscribe	Waits for completion of the Subscribe or UnSubscribe method.	
MqttClient.UnSubscribe		
MqttClient.Publish	Returns immediately to the application thread after passing the request to the MQTT client.	None.
MqttDeliveryToken.waitForCompletion	Waits for the delivery token to be returned.	Indefinite, or as set as a parameter.

## Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure “at least once” and “exactly once” delivery, and “exactly once” receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

When you connect an MQTT client application using the `MqttClient.connect` method, the client identifies the connection using the client identifier and the address of the server. The server checks whether session information has been saved from a previous connection to the server. If a previous session still exists, and `cleanSession=true`, then the previous session information at the client and server is cleared. If `cleanSession=false` the previous session is resumed. If no previous session exists, a new session is started.

**Note:** The IBM MQ Administrator can forcibly close an open session and delete all the session information. If the client reopens the session with `cleanSession=false`, a new session is started.

## Publications

If you use the default `MqttConnectOptions`, or set `MqttConnectOptions.cleanSession` to `true` before connecting the client, all pending publication deliveries for the client are removed when the client connects.

The clean session setting has no effect on publications sent with `QoS=0`. For `QoS=1` and `QoS=2`, using `cleanSession=true` might result in losing a publication.

## Subscriptions

If you use the default `MqttConnectOptions`, or set `MqttConnectOptions.cleanSession` to `true` before connecting the client, any old subscriptions for the client are removed when the client connects. Any new subscriptions the client makes during the session are removed when it disconnects.

If you set `MqttConnectOptions.cleanSession` to `false` before connecting, any subscriptions the client creates are added to all the subscriptions that existed for the client before it connected. All the subscriptions remain active when the client disconnects.

Another way of understanding the way the `cleanSession` attribute affects subscriptions is to think of it as a modal attribute. In its default mode, `cleanSession=true`, the client creates subscriptions and receives publications only within the scope of the session. In the alternative mode, `cleanSession=false`, subscriptions are durable. The client can connect and disconnect and its subscriptions remain active. When the client reconnects, it receives any undelivered publications. While it is connected, it can modify the set of subscriptions that are active on its behalf.

You must set the `cleanSession` mode before connecting; the mode lasts for the whole session. To change its setting, you must disconnect and reconnect the client. If you change modes from using `cleanSession=false` to `cleanSession=true`, all previous subscriptions for the client, and any publications that have not been received, are discarded.

## Client identifier

The client identifier is a 23 byte string that identifies an MQTT client. Each identifier must be unique to only one connected client at a time. The identifier must contain only characters valid in a queue manager name. Within these constraints, you are able to use any identification string. It is important to have a procedure for allocating client identifiers, and a means of configuring a client with its chosen identifier.

The client identifier is used in the administration of an MQTT system. With potentially hundreds of thousands of clients to administer, you need to be able to identify a particular client rapidly. Suppose for example, a device has malfunctioned, and you are notified, perhaps by a customer ringing a help desk. How does the customer identify the device, and how do you correlate that identification with the server that is typically connected to the client? Do you have to consult a database that maps each device to a client identifier and to a server? Does the name of the device identify which server it is attached to? When you browse through MQTT client connections, each connection is labeled with the client identifier. Do you need to look up a table to map a client identifier to a physical device?

Does the client identifier identify a particular device, a user, or an application running at the client? If a customer replaces a faulty device with a new one, does the new device have the same identifier as the old device? Do you allocate a new identifier? If you change a physical device, but keep the same identifier, outstanding publications and active subscriptions are automatically transferred to the new device.

How do you ensure that client identifiers are unique? As well as a system for generating unique identifiers, you must have a reliable process for setting the identifier on the client. Perhaps the client device is a "black-box", with no user interface. Do you manufacture the device with a client identifier - such as using its MAC address? Or do you have a software installation and configuration process that configures the device before it is activated?

You might create a client identifier from the 48 bit device MAC address, to keep the identifier short and unique. If transmission size is not a critical issue, you might use the remaining 17 bytes to make the address easier to administer.

## Delivery tokens

When a client publishes on a topic a new delivery token is created. Use the delivery token to monitor the delivery of a publication, or to block the client application until delivery is complete.

The token is an `MqttDeliveryToken` object. It is created by calling the `MqttTopic.publish()` method and is retained by the MQTT client until the client session is disconnected and the delivery is completed.

The normal use of the token is to check whether delivery is complete. Block the client application until delivery is complete by using the returned token to call `token.waitForCompletion`. Alternatively, provide a `MqttCallback` handler. When the MQTT client has received all the acknowledgments it expects as part of delivering the publication, it calls `MqttCallback.deliveryComplete` passing the delivery token as a parameter.

Until delivery is complete, you can inspect the publication using the returned delivery token by calling `token.getMessage`.

## Completed deliveries

The completion of deliveries is asynchronous and depends on the quality of service associated with the publication.

### At most once

`QoS=0`

Delivery is complete immediately on return from `MqttTopic.publish`. `MqttCallback.deliveryComplete` is called immediately.

### At least once

`QoS=1`

Delivery is complete when an acknowledgment to the publication has been received from the queue manager. `MqttCallback.deliveryComplete` is called when the acknowledgment is received. The message might be delivered more than once before `MqttCallback.deliveryComplete` is called, if communications are slow or unreliable.

### Exactly once

`QoS=2`

Delivery is complete when the client receives a completion message that the publication has been published to subscribers. `MqttCallback.deliveryComplete` is called as soon as the publication message is received. It does not wait for the completion message.

In rare circumstances, your client application might not return to the MQTT client from `MqttCallback.deliveryComplete` normally. You know that delivery has completed, because the `MqttCallback.deliveryComplete` was called. If the client restarts the same session, `MqttCallback.deliveryComplete` does not get called again.

## Incomplete deliveries

If the delivery is not complete after the client session is disconnected you can connect the client again and complete the delivery. You can only complete the delivery of a message if the message was published in a session with the `MqttConnectionOptions` attribute set to `false`.

Create the client using the same client identifier and server address, and then connect, setting the `cleanSession` `MqttConnectionOptions` attribute to `false` again. If you set `cleanSession` to `true`, pending delivery tokens are thrown away.

You can check if there are any pending deliveries by calling `MqttClient.getPendingDeliveryTokens`. You can call `MqttClient.getPendingDeliveryTokens` before connecting the client.



## Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Create a topic for the last will and testament. You might create a topic such as `MQTTManagement/Connections/server URI/client identifier/Lost`.

Set up a "last will and testament" using the `MqttConnectionOptions.setWill(MqttTopic lastWillTopic, byte [] lastWillPayload, int lastWillQos, boolean lastWillRetained)` method.

Consider creating a time stamp in the `lastWillPayload` message. Include other client information that assists in identifying the client and the circumstances of the connection. Pass the `MqttConnectionOptions` object to the `MqttClient` constructor.

Set `lastWillQos` to 1 or 2, to make the message persistent in IBM MQ, and to guarantee delivery. To retain the last lost connection information, set the `lastWillRetained` to `true`.

The "last will and testament" publication is sent to subscribers if the connection ends unexpectedly. It is sent if the connection ends without the client calling the `MqttClient.disconnect` method.

To monitor connections, complement the "last will and testament" publication with other publications to record connections and programmed disconnections.

## Message persistence in MQTT clients

Publication messages are made persistent if they are sent with a quality of service of "at least once", or "exactly once". You can implement your own persistence mechanism on the client, or use the default persistence mechanism that is provided with the client. Persistence works in both directions, for publications sent to or from the client.

In MQTT, message persistence has two aspects; how the message is transferred, and whether it is queued in IBM MessageSight and IBM MQ as a persistent message.

1. The MQTT client couples message persistence with quality of service. Depending on what quality of service you choose for a message, the message is made persistent. Message persistence is necessary to implement the required quality of service.

If you specify "at most once", `QoS=0`, the client discards the message as soon as it is published. If there is any failure in the upstream processing of the message, the message is not sent again. Even if the client remains active the message is not sent again. The behavior of `QoS=0` messages is the same as IBM MQ fast nonpersistent messages.

If a message is published by a client with `QoS` of 1 or 2, it is made persistent. The message is stored locally, and only discarded from the client when it is no longer needed to guarantee "at least once", `QoS=1`, or "exactly once", `QoS=2`, delivery.

2. If a message is marked as `QoS` 1 or 2, it is queued in IBM MessageSight and IBM MQ as a persistent message. If it is marked as `QoS=0`, then it is queued in IBM MessageSight and IBM MQ as a nonpersistent message. In IBM MQ nonpersistent messages are transferred between queue managers "exactly once", unless the message channel has the `NPMSPEED` attribute set to `FAST`.

A persistent publication is stored on the client until it is received by a client application. For `QoS=2`, the publication is discarded from the client when the application callback returns control. For `QoS=1` the application might receive the publication again, if a failure occurs. For `QoS=0`, the callback receives the publication no more than once. It might not receive the publication if there is a failure, or if the client is disconnected at the time of publication.

When you subscribe to a topic, you can reduce the QoS with which the subscriber receives messages to match its persistence capabilities. Publications that are created at a higher QoS are sent with the highest QoS that the subscriber requested.

## Storing messages

The implementation of data storage on small devices varies a great deal. The model of temporarily saving persistent messages in storage that is managed by the MQTT client might be too slow, or demand too much storage. In mobile devices, the mobile operating system might provide a storage service that is ideal for MQTT messages.

To provide flexibility in meeting the constraints of small devices, the MQTT client has two persistence interfaces. The interfaces define the operations that are involved in storing persistent messages. The interfaces are described in the API documentation for the MQTT client for Java. For links to client API documentation for the MQTT client libraries, see MQTT client programming reference. You can implement the interfaces to suit a device. The MQTT client that runs on Java SE has a default implementation of the interfaces that store persistent messages in the file system. It uses the `java.io` package.

## Persistence classes

### MqttClientPersistence

Pass an instance of your implementation of `MqttClientPersistence` to the MQTT client as a parameter of the `MqttClient` constructor. If you omit the `MqttClientPersistence` parameter from the `MqttClient` constructor, the MQTT client stores persistent messages using the class `MqttDefaultFilePersistence`.

### MqttPersistable

`MqttClientPersistence` gets and puts `MqttPersistable` objects using a storage key. You must provide an implementation of `MqttPersistable` as well as the implementation of `MqttClientPersistence` if you are not using the `MqttDefaultFilePersistence`.

### MqttDefaultFilePersistence

The MQTT client provides the `MqttDefaultFilePersistence` class. If you instantiate `MqttDefaultFilePersistence` in your client application, you can provide the directory to store persistent messages as a parameter of the `MqttDefaultFilePersistence` constructor.

Alternatively, the MQTT client can instantiate `MqttDefaultFilePersistence` and place files in a default directory. The name of the directory is *client identifier -tcp hostname portnumber*. `"\"`, `"\\"`, `"/"`, `":"` and `" "` are removed from the directory name string.

The path to the directory is the value of the system property `rcp.data`. If `rcp.data` is not set, the path is the value of the system property `usr.data`.

`rcp.data` is a property associated with installation of an OSGi or Eclipse Rich Client Platform (RCP).

`usr.data` is the directory in which the Java command that started the application was launched.

## Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT clients can create publications to send to IBM MQ, and subscribe to topics on IBM MQ to receive publications.

An `MqttMessage` has a byte array as its payload. Aim to keep messages as small as possible. The maximum length of message permitted by the MQTT protocol is 250 MB.

Typically, an MQTT client program uses `java.lang.String` or `java.lang.StringBuffer` to manipulate message contents. For convenience, the `MqttMessage` class has a `toString` method to convert its payload to a string. To create the byte array payload from a `java.lang.String` or `java.lang.StringBuffer`, use the `getBytes` method.

The `getBytes` method converts a string to the default character set for the platform. The default character set is generally UTF-8. MQTT publications that contain only text are usually encoded in UTF-8. Use the method `getBytes("UTF8")` to override the default character set.

In IBM MQ, an MQTT publication is received as a `jms-bytes` message. The message includes an `MQRFH2` folder containing an `<mqtt>`, and an `<myps>` folder. The `<mqtt>` folder contains the `clientId`, `msgId`, and `qos`, but this content might change in the future.

An `MqttMessage` has three additional attributes: quality of service, whether it is retained, and whether it is a duplicate. The duplicate flag is only set if the quality of service is "at least once" or "exactly once". If the message was sent previously, and not acknowledged quickly enough by the MQTT client, the message is sent again, with the duplicate attribute set to `true`.

## Publishing

To create a publication in an MQTT client application, create an `MqttMessage`. Set its payload, quality of service and whether it is retained, and call the `MqttTopic.publish(MqttMessage message)` method; `MqttDeliveryToken` is returned and the completion of the publication is asynchronous.

Alternatively, the MQTT client can create a temporary message object for you from the parameters on the `MqttTopic.publish(byte [] payload, int qos, boolean retained)` method when it creates a publication.

If the publication has an "at least once" or an "exactly once" quality of service, `QoS=1` or `QoS=2`, the MQTT client calls the `MqttClientPersistence` interface. It calls `MqttClientPersistence` to store the message before returning a delivery token to the application.

The application can choose to block until the message is delivered to the server, using the `MqttDeliveryToken.waitForCompletion` method. Alternatively, the application can continue without blocking. If you want to check if publications are delivered, without blocking, register an instance of a callback class that implements `MqttCallback` with the MQTT client. The MQTT client calls the `MqttCallback.deliveryComplete` method as soon as the publication has been delivered. Depending on the quality of service, the delivery might be almost immediate for `QoS=0`, or it might take some time for `QoS=2`.

Use the `MqttDeliveryToken.isComplete` method to poll if delivery is complete. While the value of `MqttDeliveryToken.isComplete` is `false`, you can call `MqttDeliveryToken.getMessage` to get the message contents. If the result of calling `MqttDeliveryToken.isComplete` is `true`, the message has been discarded and calling `MqttDeliveryToken.getMessage` would throw a null pointer exception. There is no built-in synchronization between `MqttDeliveryToken.getMessage` and `MqttDeliveryToken.isComplete`.

If the client disconnects before receiving all the pending delivery tokens, a new instance of the client can query pending delivery tokens before connecting. Until the client connects, no new deliveries are completed, and it is safe to call `MqttDeliveryToken.getMessage`. Use the `MqttDeliveryToken.getMessage` method to find out which publications have not been delivered. Pending delivery tokens are discarded if

you connect with `MqttConnectOptions.cleanSession` set to its default value, `true`.

## Subscribing

A queue manager or IBM MessageSight is responsible for creating publications to send to an MQTT subscriber. The queue manager checks if the topic filter in a subscription created by an MQTT client matches the topic string in a publication. The match can either be an exact match, or the match can include wildcards. Before the publication is forwarded to the subscriber by the queue manager, the queue manager checks the topic attributes associated with the publication. It follows the search procedure described in *Subscribing using a topic string that contains wildcard characters to identify if an administrative topic object grants the user authority to subscribe*.

When the MQTT client receives a publication with "at least once" quality of service, it calls the `MqttCallback.messageArrived` method to process the publication. If the quality of service of the publication is "exactly once", `QoS=2`, the MQTT client calls the `MqttClientPersistence` interface to store the message when it is received. It then calls `MqttCallback.messageArrived`.

## Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to IBM MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to IBM MQ to create a subscription, the request is sent with the "at least once" quality of service.

The quality of service of a publication is an attribute of `MqttMessage`. It is set by the method `MqttMessage.setQos`.

The method `MqttClient.subscribe` can reduce the quality of service applied to publications sent to a client on a topic. The quality of service of a publication forwarded to a subscriber might be different to the quality of service of the publication. The lower of the two values is used to forward a publication.

### At most once

`QoS=0`

The message is delivered at most once, or it is not delivered at all. Its delivery across the network is not acknowledged.

The message is not stored. The message might be lost if the client is disconnected, or if the server fails.

`QoS=0` is the fastest mode of transfer. It is sometimes called "fire and forget".

The MQTT protocol does not require servers to forward publications at `QoS=0` to a client. If the client is disconnected at the time the server receives the publication, the publication might be discarded, depending on the server. The telemetry (MQXR) service does not discard messages sent with `QoS=0`. They are stored as nonpersistent messages, and are only discarded if the queue manager stops.

### At least once

`QoS=1`

`QoS=1` is the default mode of transfer.

The message is always delivered at least once. If the sender does not receive an acknowledgment, the message is sent again with the DUP flag set until an acknowledgment is received. As a result, the receiver can be sent the same message multiple times, and might process it multiple times.

The message must be stored locally at the sender and the receiver until it is processed.

The message is deleted from the receiver after it has processed the message. If the receiver is a broker, the message is published to its subscribers. If the receiver is a client, the message is delivered to the subscriber application. After the message is deleted, the receiver sends an acknowledgment to the sender.

The message is deleted from the sender after it has received an acknowledgment from the receiver.

### **Exactly once**

QoS=2

The message is always delivered exactly once.

The message must be stored locally at the sender and the receiver until it is processed.

QoS=2 is the safest, but slowest mode of transfer. It takes at least two pairs of transmissions between the sender and receiver before the message is deleted from the sender. The message can be processed at the receiver after the first transmission.

In the first pair of transmissions, the sender transmits the message and gets acknowledgment from the receiver that it has stored the message. If the sender does not receive an acknowledgment, the message is sent again with the DUP flag set until an acknowledgment is received.

In the second pair of transmissions, the sender tells the receiver that it can complete processing the message, "PUBREL". If the sender does not receive an acknowledgment of the "PUBREL" message, the "PUBREL" message is sent again until an acknowledgment is received. The sender deletes the message it saved when it receives the acknowledgment to the "PUBREL" message

The receiver can process the message in the first or second phases, provided that it does not reprocess the message. If the receiver is a broker, it publishes the message to subscribers. If the receiver is a client, it delivers the message to the subscriber application. The receiver sends a completion message back to the sender that it has finished processing the message.

### **Retained publications and MQTT clients**

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Use the `MqttMessage.setRetained` method to specify whether a publication on a topic is retained, or not.

To delete a retained publication in IBM MQ, run the `CLEAR TOPICSTR CLEAR TOPICSTR MQSC` command.

If you create a publication with a null payload, the empty publication is forwarded to subscribers. Other MQTT brokers might not forward an empty publication to subscribers.

If you publish a non-retained publication to a topic that has a retained publication, the retained publication is not affected. Current subscribers receive the new publication. New subscribers receive the retained publication first, then receive any new publications.

When you create or update a retained publication, send the publication with a QoS of 1 or 2. If you send it with a QoS of 0, IBM MQ creates a nonpersistent retained publication. The publication is not retained if the queue manager stops.

Use retained publications to record the latest value of a measurement. New subscribers to the retained topic immediately receive the most recent value of the measurement. If no new measurements are taken since the subscriber last subscribed to the publication topic, and if the subscriber subscribes again, the subscriber receives the most recent retained publication on the topic again.

## Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Create subscriptions using the `MqttClient.subscribe` methods, passing one or more topic filters and quality of service parameters. The quality of service parameter sets the maximum quality of service that the subscriber is prepared to use to receive a message. Messages sent to this client cannot be delivered with a higher quality of service. The quality of service is set to the lower of the original value when the message was published and the level specified for the subscription. The default quality of service for receiving messages is `QoS=1`, at least once.

The subscription request itself is sent with `QoS=1`.

Publications are received by a subscriber when the MQTT client calls the `MqttCallback.messageArrived` method. The `messageArrived` method also passes the topic string with which the message was published to the subscriber.

You can remove a subscription, or a set of subscriptions, using the `MqttClient.unsubscribe` methods.

An IBM MQ command can remove a subscription. List subscriptions using IBM MQ Explorer, or by using `runmqsc` or PCF commands. All MQTT client subscriptions are named. They are given a name of the form: *ClientIdentifier:Topic name*

If you use the default `MqttConnectOptions`, or set `MqttConnectOptions.cleanSession` to `true` before connecting the client, any old subscriptions for the client are removed when the client connects. Any new subscriptions the client makes during the session are removed when it disconnects.

If you set `MqttConnectOptions.cleanSession` to `false` before connecting, any subscriptions the client creates are added to all the subscriptions that existed for the client before it connected. All the subscriptions remain active when the client disconnects.

Another way of understanding the way the `cleanSession` attribute affects subscriptions is to think of it as a modal attribute. In its default mode, `cleanSession=true`, the client creates subscriptions and receives publications only within the scope of the session. In the alternative mode, `cleanSession=false`, subscriptions are durable. The client can connect and disconnect and its subscriptions remain active. When the client reconnects, it receives any undelivered publications. While it is connected, it can modify the set of subscriptions that are active on its behalf.

You must set the `cleanSession` mode before connecting; the mode lasts for the whole session. To change its setting, you must disconnect and reconnect the client. If you change modes from using `cleanSession=false` to `cleanSession=true`, all previous subscriptions for the client, and any publications that have not been received, are discarded.

Publications that match active subscriptions are sent to the client as soon as they are published. If the client is disconnected, they are sent to the client if it reconnects to the same server with the same client identifier and `MqttConnectOptions.cleanSession` set to `false`.

Subscriptions for a particular client are identified by the client identifier. You can reconnect the client from a different client device to the same server, and continue with the same subscriptions and receive undelivered publications.

## Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM MQ.

Topics strings are used to send publications to subscribers. Create a topic string using the method, `MqttClient.getTopic(java.lang.String topicString)`.

Topic filters are used to subscribe to topics and receive publications. Topic filters can contain wildcards. With wildcards, you can subscribe to multiple topics. Create a topic filter by using a subscription method; for example, `MqttClient.subscribe(java.lang.String topicFilter)`.

### Topic strings

The syntax of an IBM MQ topic string is described in [Topic Strings](#). The syntax of MQTT topic strings is described in the `MqttClient` class in the API documentation for the MQTT client for Java. For links to client API documentation for the MQTT client libraries, see [MQTT client programming reference](#).

The syntax of each type of topic string is almost identical. There are four minor differences:

1. Topic strings sent to IBM MQ by MQTT clients must follow the convention for queue manager names.
2. The maximum lengths differ. IBM MQ topic strings are limited to 10,240 characters. An MQTT client can create topic strings of up to 65535 bytes.
3. A topic string created by an MQTT client cannot contain a null character.
4. In IBM Integration Bus, a null topic level, '`...//...`' is invalid. Null topic levels are supported by IBM MQ.

Unlike IBM MQ `publish/subscribe`, the `mqttv3` protocol does not have a concept of an administrative topic object. You cannot construct a topic string from a topic object and a topic string. However, a topic string is mapped to an administrative topic in IBM MQ. The access control associated with the administrative topic determines whether a publication is published to the topic, or discarded. The attributes that are applied to a publication when it is forwarded to subscribers, are influenced by the attributes of the administrative topic.

### Topic filters

The syntax of an IBM MQ topic filter is described in [Topic-based wildcard scheme](#). The syntax of the topic filters you can construct with an MQTT client are described in the `MqttClient` class in the API documentation for the MQTT client for Java. For links to client API documentation for the MQTT client libraries, see [MQTT client programming reference](#).

## C client programming concepts

Links to the latest differences between the C and Java client for the IBM MQ Telemetry Transport are given in this topic. The topic complements the client concepts and the C reference information.

For further information, see the:

- Eclipse Paho web site, for the latest telemetry clients and an Android sample.
- MA9C: IBM Messaging Telemetry Clients SupportPac, which includes a similar set of clients to Paho, plus samples for Java and C.
- Deprecated MA9B: Mobile Messaging and M2M Client Pack includes a sample for iOS that was written in the IBM WebSphere MQ Version 7.5 time frame.

---

## Developing MQ Light applications

> V9.0.0

IBM MQ Light is a product and messaging API, aimed at making it quick and easy for developers to make their applications more responsive and scalable.

The IBM MQ support for MQ Light APIs allows an IBM MQ administrator to define a new type of channel: an AMQP channel. When the AMQP channel is started it defines a port number, which accepts connections from MQ Light applications.

The MQ Light API is based on the Oasis AMQP 1.0 protocol. There are messaging APIs for Node.js, Java, Ruby and Python, available at <https://github.com/mqlight>.

An application that is developed to use the MQ Light API can be connected to either an MQ Light runtime, an IBM MQ queue manager with an AMQP channel, or an instance of an MQ Light service in IBM Cloud (formerly Bluemix®).

### Developing applications

MQ Light aims to make it easier to prototype and develop business applications rapidly.

When an application is ready to deploy, it requires all of the monitoring, reliability, and security capabilities of other enterprise applications. It can also exchange data with other enterprise applications. You can deploy MQ Light applications to an IBM MQ queue manager.

When you have installed and connected an AMQP client, you can exchange messages with IBM MQ applications. For example, if you use the MQ Light Node.js client to send a JavaScript string message, the IBM MQ application receives an MQ message, where the format field of the MQMD is set to MQSTR.

For information about securing MQ Light applications, see MQ Light security.

### Managing the AMQP channel

The AMQP channel can be managed in the same way as other MQ channels. You can use MQSC commands, PCF command messages, or IBM MQ Explorer to define, start, stop, and manage the channels. In *Creating and using AMQP channels*, example commands are provided to define and start connecting clients to a queue manager.

When an AMQP channel is started, you can test it by connecting an MQ Light application, by using any of the following methods:

- Using the IBM MQ Light client for Node.js and Java.
- Using the IBM MQ Light early access program client for Ruby and Python.



- Using another AMQP Version 1.0 client. For example, Apache Qpid Proton.

## Trying out MQ Light

The MQ Light client downloads all include several samples, which demonstrate the different messaging features:

- Send sample
- Receive sample
- UI Workout sample

## Downloading a client

IBM MQ does not ship MQ Light clients but you can download and install the following MQ Light clients:

### Node.js

Install MQ Light Node.js API to your working directory using npm: `npm install mqlight@1.0`

**Java** Download the `mqlight-distribution` package for the required version from Maven Central and extract the contents. You can find the available versions of the `mqlight-distribution` packages on Maven Central.

**Ruby** Install MQ Light Ruby API to your working directory using gem: `gem install mqlight --pre`

### Python

Install MQ Light Python API to your working directory using pip: `pip install mqlight --pre`

For more information, see Obtaining an MQ Light client.

You can also download other open-source AMQP clients based on Apache Qpid libraries, For more information, see <https://qpid.apache.org/index.html>

## Finding out more about MQ Light

You can find out more information about MQ Light and also connect with developers and other users at the DevWorks community website, which includes blog articles and a forum.

## AMQP and IBM MQ application interaction



If you are developing a system that incorporates AMQP applications, you need to understand how AMQP and IBM MQ applications can interact.

You need to understand how AMQP message formats relate to IBM MQ message formats, including how AMQP properties and AMQP application properties are mapped to IBM MQ MD fields and IBM MQ message properties. The following topics describe the AMQP message format, and how it maps onto IBM MQ message format.

## IBM MQ Light quality of service with AMQP

V9.0.0

There are four features of the IBM MQ Light API that allow you to control the reliability of message delivery to, and from, MQ Light and AMQP applications.

These are:

- “Message quality of service (QOS)”
- “Subscriber auto-confirm” on page 1171
- “Subscription time to live” on page 1171
- “Message persistence” on page 1171

### Message quality of service (QOS)

The MQ Light API offers two qualities of service:

- At most once
- At least once

You can choose which quality of service you want publishers and subscribers to use.

If you are using an MQ Light client, set the client or subscribe **qos** option to *QOS\_AT\_MOST\_ONCE* or *QOS\_AT\_LEAST\_ONCE*.

If you are using a different AMQP client, set the **settled** attribute of the transfer frame (for publishers), or the disposition frame (for subscribers) to *true* or *false*, depending on the quality of service you want to achieve.

The quality of service determines when a message is discarded from the sending side of a conversation.

#### Publishing

If a publisher chooses **QOS 0** (at most once) the publisher does not wait for an acknowledgment from the queue manager, before it discards its copy of the message.

If the connection to the queue manager fails before the send has completed, the message might not be received by subscribers.

If a publisher chooses **QOS 1** (at least once) the publisher waits for the queue manager to acknowledge that the message has been written to subscriber queues before discarding its copy of the message.

If the connection to the queue manager fails during the send, the publisher re-sends the message once it has reconnected to the queue manager.

#### Subscribing

If a subscriber chooses **QOS 0** the queue manager does not wait for an acknowledgment from the subscriber before discarding its copy of the message.

If the connection to the subscriber fails before the subscriber has received the message, that message might be lost.

If a subscriber chooses **QOS 1** the queue manager waits for an acknowledgment from the subscriber before discarding its copy of the message.

If the connection to the subscriber fails before the subscriber has received the message, the message is kept by the queue manager. The queue manager re-sends the message to the subscriber when the queue manager reconnects, or to another subscriber if the subscription is shared.

### Subscriber auto-confirm

If a subscriber chooses **QOS 1** (at least once) it must acknowledge receipt of each message before the queue manager discards its copy. The subscriber can decide when it acknowledges messages.

With **auto-confirm** set to *true*, the MQ Light client automatically acknowledges delivery of each message, once the client has successfully received the message over the network.

This ensures that if there is a network failure, the message is redelivered to the application. However, it is still possible for the application to lose the message, if the application fails in between the MQ Light client acknowledging the message, and the application processing it.

With **auto-confirm** set to *false*, the MQ Light client does not automatically acknowledge delivery of the message, but leaves it to the application to decide when it should be acknowledged.

This allows an application to make an update to an external resource, such as a database or a file, before acknowledging to the queue manager that the message has now been processed and can be discarded.

### Subscription time to live

When an application subscribes, it chooses whether the subscription, and the destination where messages are stored for that subscription, continue to exist after the application disconnects.

The MQ Light subscribe option **ttl** is used to specify the time (in milliseconds) that a subscription continues to exist after the application disconnects. If the application reconnects before this time, the subscription is resumed and the application can continue to consume messages from that subscription.

If the time-to-live period passes without the application reconnecting, the subscription is removed and any messages stored on its destination are lost, even if they are persistent messages.

If it is important not to lose messages, you must specify a time-to-live value to the application, that is high enough to ensure messages are not lost during an outage.

### Message persistence

The persistence of messages is controlled by the publishing and subscribing applications, and the configuration of IBM MQ topic objects.

If the AMQP subscriber uses **QOS 0** (at most once) and creates a non-durable subscription, the AMQP channel always puts non-persistent messages to the subscriber queue, regardless of the other options described in the following text.

Note that, if the queue manager is stopped both the subscription and the messages will be lost.

If an AMQP publisher sets the AMQP  **durable**  header to *true*, the AMQP channel puts persistent messages to subscriber queues.

If the queue manager is stopped for any reason, the messages are still available to subscribers when the queue manager is restarted.

If the  **durable**  header is not set, the AMQP channel chooses the persistence of published messages based on the  **DEFPSIST**  attribute of the relevant IBM MQ topic object.

By default, this is the SYSTEM.BASE.TOPIC, which uses a **DEFPSIST** attribute of *NO* (non-persistent).

## AMQP messages

V 9.0.0

AMQP messages are composed of a header, delivery annotations, message annotations, properties, application properties, body, and footer.

AMQP messages are composed of the following parts:

### Header

The optional header contains five fixed attributes of the message:

- **durable**  - specifies durability requirements
- **priority**  - relative message priority
- **ttl**  - time to live in milliseconds
- **first-acquirer**  - if this is true, the message has not been acquired by any other link
- **delivery-count**  - the number of previous, unsuccessful delivery attempts.

### Delivery-annotations

Optional. Specifies non-standard header attributes of the message for different intended audiences. Delivery annotations convey information from the sending peer to the receiving peer.

### Message-annotations

Optional. Specifies non-standard header attributes of the message for different intended audiences. The message-annotations section is used for properties of the message which are aimed at the infrastructure and should be propagated across every delivery step.

### Properties

Optional. This part is equivalent the MQ message descriptor. It contains the following fixed fields:

- **message-id**  - application message identifier
- **user-id**  - ID of creating user
- **to**  - address of node that the message is destined for
- **subject**  - the subject of the message
- **reply-to**  - the node that the send replies to
- **correlation-id**  - application correlation identifier
- **content-type**  - MIME content type
- **content-encoding**  - MIME content type. Used as a modifier to the content-type.
- **absolute-expiry-time**  - the time when this message is considered expired
- **creation-time**  - the time when this message was created
- **group-id**  - the group that this message belongs to
- **group-sequence**  - the sequence number of this message within its group
- **reply-to-group-id**  - the group that the reply message belongs to


### Applications-properties

Equivalent to MQ message properties.

**Body** Equivalent to the MQ user payload.

**Footer** Optional. The footer is used for details about the message or delivery that can only be calculated or evaluated after the whole bare message has been constructed or seen (for example, message hashes, HMACs, signatures and encryption details).

The AMQP message format is illustrated in the following figure:



The properties, application-properties, and application-data part are known as the “bare message”. This is the message as sent by the sender, and is immutable. The receiver sees the entire message, including the header, footer, delivery-annotations and message-annotations.

For a full description of the AMQP 1.0 message format, see the OASIS Standard at <http://docs.oasis-open.org/amqp/core/v1.0/amqp-core-complete-v1.0.pdf>.

### Mapping IBM MQ fields onto AMQP fields (outgoing messages): V 9.0.0

When an IBM MQ message is published and IBM MQ sends it to an AMQP consumer, it will propagate some of the attributes of the IBM MQ message into equivalent AMQP message attributes.

#### header

A header is only included if one of the five fields in the header contains a non-default value. Only the fields with a non-default value are included in the header. The five header fields are initially derived from the equivalent `mq_amqp.Hdr` property, if it is set, and then modified as shown in the following table:

Table 166. Header field mappings

Field	Default value	Value
durable	false	True if <code>MQMD.Persistence</code> is set to <code>MQPER_PERSISTENT</code> , false otherwise.
priority	4	From <code>mq_amqp.Hdr.Pri</code> , if set, or otherwise from <code>MQMD.Priority</code> , if set. If neither set, set to 4.
ttl	n/a	<code>MQMD.Expiry</code> in milliseconds. If the value of <code>MQMD.Expiry</code> is <code>MQEI_UNLIMITED</code> , then set to the maximum value for the AMQP ttl field
first-acquirer	false	From <code>mq_amqp.Hdr.Fac</code> , if set, or false otherwise.
delivery-count	0	From <code>mq_amqp.Hdr.Dct</code> , if set, or 0 otherwise.

#### delivery-annotation

Set as necessary by the AMQP channel.

#### message-annotation

Not included.

#### properties

The **properties** will come unmodified from the equivalent `mq_amqp.Prp` properties if these are set. If the message was not originally an AMQP message (that is, `PutApplType` is not `MQAT_AMQP`), then a properties section is generated as described in the following table:

Table 167. Properties field mappings

Name	Value
message-id	The MQMD.MsgId is set as binary.
user-id	The UTF-8 form of the MQMD.UserIdentifier is set as binary in network byte-order.
to	The queue that the message was got from, or, for a publication, the topic string.
subject	Not set.
reply-to	The MQMD.ReplyToQ if non-blank, otherwise not set.
correlation-id	The MQMD.CorrelId is set as binary if non-blank, otherwise not set.
content-type	Not set.
content-encoding	Not set.
absolute-expiry-time	Not set.
creation-time	The MQMD.PutDate and MQMD.PutTime fields are used to generate a timestamp.
group-id	Not set.
group-sequence	Not set.
reply-to-group-id	Not set.

### application-properties

All IBM MQ properties in the “usr” group are added as the **application-properties**.

### body

The AMQP channel performs a get with convert, to convert the IBM MQ payload into UTF-8.

If the IBM MQ payload does not contain an AMQP message, then the IBM MQ payload is set in the body as a single string data section for Format MQFMT\_STRING (provided conversion to UTF-8 was successful), or as a single binary data section otherwise.

If an AMQP format message is included, then this is set as the body. Any IBM MQ headers (not including the messages properties which are returned in a message handle) that precede the AMQP message are prepended as a binary value if the body is an AMQP Sequence. Otherwise the IBM MQ headers are discarded.

### footer

No footer is included.

## Related information:

MQMD - Message descriptor

## Mapping AMQP fields onto IBM MQ fields (incoming messages): V9.0.0

When the AMQP channel receives a message and puts it to IBM MQ, it propagates some of the attributes of the AMQP message into equivalent IBM MQ message attributes.

The following restrictions apply when mapping an incoming AMQP message:

- If the **message-id** or **correlation-id** field in the properties part is a uuid or a ulong, then the message is rejected.
- Any **message-annotations** cause the message to be rejected.
- **delivery-annotations** and **footer** sections are allowed, but are not propagated into the IBM MQ message.

The following sub-sections show the IBM MQ expression of an AMQP message.

### message descriptor

Table 168. Message descriptor for AMQP message

Field	Value
StrucId	MQMD_STRUC_ID
Version	MQMD_VERSION_1
Report	MQRO_NONE
MsgType	MQMT_DATAGRAM
Expiry	Value taken from the <b>ttl</b> field in the AMQP message header
Feedback	MQFB_NONE
Encoding	MQENC_NORMAL
CodedCharSetId	1208 (UTF-8)
Format	See Payload
Priority	Value taken from the <b>priority</b> field in the AMQP message header. If set, limited to a maximum of 9. If not set, takes the default value of 4.
Persistence	If the <b>durable</b> field in the AMQP message header is set true, set to MQPER_PERSISTENT. Otherwise, set to MQPER_NOT_PERSISTENT.
MagId	The queue manager allocates a unique 24-byte MsgId.
CorrelId	Value taken from the <b>correlation-id</b> field in the AMQP properties, if set. Set to a 24-byte binary value. Otherwise, set to MQCL_NONE/.
BackoutCount	0
ReplyToQ	""
ReplyToQMgr	""
UserIdentifier	Set to the identifier of the authenticated user that connected to the AMQP channel
AccountingToken	MQACT_NONE
ApplIdentityData	Hexadecimal string. Set to the last 8 bytes of the MQ connection identifier of the AMQP channel.
PutAppType	MQAT_AMQP
PutAppName	

Table 168. Message descriptor for AMQP message (continued)

Field	Value
PutDate	Value taken from the <b>creation-time</b> field of the AMQP properties, if set. Otherwise set to the current date.
PutTime	Value taken from the <b>creation-time</b> field of the AMQP properties, if set. Otherwise set to the current time.
ApplOriginData	""

### message properties

There are two reasons for setting message properties:

- To allow parts of the AMQP message to flow through the queue manager without affecting the payload of the message.
- To allow selection of the **application-properties**.

The following table shows the properties that are set from the AMQP message:

Table 169.

Property name	MQRFH2 name	Type	Description
AMQPListener	mq_amqp.Lis	MQTYPE_STRING	An identifying string for the AMQP channel. It is used to generate the message, so that interested parties can tell which version put the message (for example, the service team when diagnosing problems). The value is not validated by the queue manager, and must not be documented externally.
AMQPVersion	mq_amqp.Ver	MQTYPE_STRING	The version of the AMQP message. If not present, "1.0" is assumed. The value is not validated by the queue manager.
AMQPClient	mq_amqp.Cli	MQTYPE_STRING	An identifying string for the API. It is used to send the AMQP message to the channel, so that interested parties can tell which version put the message (for example, the service team when diagnosing problems). The value is not validated by the queue manager and must not be documented externally.
AMQPDurable	mq_amqp.Hdr.Dur	MQTYPE_BOOLEAN	The value of the <b> durable </b> field in the AMQP message header, if set.
AMQPPriority	mq_amqp.Hdr.Pri	MQTYPE_INT32	The value of the <b> priority </b> field in the AMQP message header, if set.
AMQPTtl	mq_amqp.Hdr.Ttl	MQTYPE_INT64	The value of the <b> ttl </b> field in the AMQP message header, if set.
AMQPFIRSTAcquirer	mq_amqp.Hdr.Fac	MQTYPE_BOOLEAN	The value of the <b> first-acquirer </b> field in the AMQP message header, if set.
AMQPDeliveryCount	mq_amqp.Hdr.Dct	MQTYPE_INT64	The value of the <b> delivery-count </b> field in the AMQP message header, if set.
AMQPMsgId	mq_amqp.Prp.Mid	MQTYPE_STRING	The value of the <b> message-id </b> field in the AMQP properties, if set as a string.
		MQTYPE_BYTE_STRING	The value of the <b> message-id </b> field in the AMQP properties, if set as a byte string.



Table 169. (continued)

Property name	MQRFH2 name	Type	Description
AMQPUserId	mq_amqp.Prp.Uid	MQTYPE_BYTE_STRING	The value of the <b>user-id</b> field in the AMQP properties, if set.
AMQPTo	mq_amqp.Prp.To	MQTYPE_STRING	The value of the <b>to</b> field in the AMQP properties, if set.
AMQPSubject	mq_amqp.Prp.Sub	MQTYPE_STRING	The value of the <b>subject</b> field in the AMQP properties, if set.
AMQPReplyTo	mq_amqp.Prp.Rto	MQTYPE_STRING	The value of the <b>reply-to</b> field in the AMQP properties, if set.
AMQPCorrelationId	mq_amqp.Prp.Cid	MQTYPE_STRING	The value of the <b>correlation-id</b> field in the AMQP properties, if set as a string..
		MQTYPE_BYTE_STRING	The value of the <b>correlation-id</b> field in the AMQP properties, if set as a byte string.
AMQPContentType	mq_amqp.Prp.Cnt	MQTYPE_STRING	The value of the <b>content-type</b> field in the AMQP properties, if set.
AMQPContentEncoding	mq_amqp.Prp.Cne	MQTYPE_STRING	The value of the <b>content-encoding</b> field in the AMQP properties, if set.
AMQPAbsoluteExpiryTime	mq_amqp.Prp.Aet	MQTYPE_STRING	The value of the <b>absolute-expiry-time</b> field in the AMQP properties, if set.
AMQPCreationTime	mq_amqp.Prp.Crt	MQTYPE_STRING	The value of the <b>creation-time</b> field in the AMQP properties, if set.
AMQPGroupId	mq_amqp.Prp.Gid	MQTYPE_STRING	The value of the <b>group-id</b> field in the AMQP properties, if set.
AMQPGroupSequence	mq_amqp.Prp.Gsq	MQTYPE_INT64	The value of the <b>group-sequence</b> field in the AMQP properties, if set.
AMQPReplyToGroupId	mq_amqp.Prp.Rtg	MQTYPE_STRING	The value of the <b>reply-to-group-id</b> field in the AMQP properties, if set.

Each of the application-properties from the AMQP message is set as an IBM MQ message property. The **application-properties** section must be reconstituted identically byte-for-byte, and so the following restrictions apply:

- If an application property is rejected by the MQSETMP validation code, the message to be rejected. For example:
  - The property name is limited in length to MQ\_MAX\_PROPERTY\_NAME\_LENGTH.
  - The property name must follow the rules defined by the Java Language Specification for Java Identifiers.
  - The property name must not begin "JMS" or "usr.JMS" except for the documented JMS properties that can be set.
  - The property name must not be a SQL keyword.
- An application property containing Unicode character U+002E (".") causes the message to be rejected. The property must be expressible in the "usr" group of properties used by JMS.
- Only null, boolean, byte, short, int, long, float, double, binary and string properties are supported. An application property with any other type will cause the message to be rejected.

#### payload

- For an AMQP **body** with a single binary data section, the binary data (excluding the AMQP bits) is put as the IBM MQ payload, with a Format of MQFMT\_NONE.

- For an AMQP **body** with a single string data section, the string data (excluding the AMQP bits) is put as the IBM MQ payload, with a Format of MQFMT\_STRING.
- Otherwise, the AMQP **body** forms the payload as-is, with a Format of MQFMT\_AMQP.

---

## Developing Microsoft Windows Communication Foundation (WCF) applications with IBM MQ

The Microsoft Windows Communication Foundation (WCF) custom channel for IBM MQ sends and receives messages between WCF clients and services.

### Related concepts:

“Introduction to the use of IBM MQ custom channel for WCF with .NET 3”

Overview of the information available for programmers using the IBM MQ custom channel for Windows Communication Foundation (WCF) with .NET 3.

“Using IBM MQ custom channels for WCF” on page 1184

Overview of the information available for programmers using IBM MQ custom channels for Windows Communication Foundation (WCF).

“Using the WCF samples” on page 1204

The Windows Communication Foundation (WCF) samples provide some simple examples of how the IBM MQ custom channel can be used.

“Problem determination on the WCF custom channel for IBM MQ” on page 1210

You can use IBM MQ trace to collect detailed information about what various parts of the IBM MQ code is doing. When using Windows Communication Foundation (WCF), a separate trace output is generated for the WCF custom channel trace integrated with the Microsoft WCF infrastructure trace.

## Introduction to the use of IBM MQ custom channel for WCF with .NET 3

Overview of the information available for programmers using the IBM MQ custom channel for Windows Communication Foundation (WCF) with .NET 3.

### What is the IBM MQ custom channel for WCF?

The custom channel for IBM MQ is a transport channel using the Microsoft Windows Communication Foundation (WCF) unified programming model.

The Microsoft Windows Communication Foundation framework, introduced in Microsoft.NET 3, enables .NET applications and services to be developed independently from the transport and protocols used to connect them, enabling alternative transports or configurations to be used according to the environment that the service or application is deployed in.

Connections are managed at run time by WCF by building a channel stack containing the required combination of:

- Protocol elements: An optional set of elements where none, one, or more can be added to support protocols such as the WS-\* standards.
- Message encoder: A mandatory element in the stack controlling the serializing of the message into its wire format.
- Transport channel: A mandatory element in the stack responsible for transporting the serialized message to its endpoint.

The custom channel for IBM MQ is a transport channel, and as such must be paired with a message encoder and optional protocols as required by the application using a WCF custom binding. In this way, applications which have been developed to use WCF can use the custom channel for IBM MQ to send and receive data in the same way as they use the built-in transports provided by Microsoft, enabling simple integration with the asynchronous, scalable, and reliable messaging functions of IBM MQ. For a full list of supported functions, see: “WCF Custom channel features and capabilities” on page 1185.

## When and why do I use the IBM MQ custom channel for WCF?

You can use the IBM MQ custom channel to send and receive messages between WCF clients and services in the same way as the built-in transports provided by Microsoft, enabling applications to access the features of IBM MQ within the WCF unified programming model.


Typical usage pattern scenarios for the IBM MQ custom channel for WCF are:

- As an interface to web services hosted over IBM MQ (SOAP over JMS).
- As a non-SOAP interface for transmission of native IBM MQ messages.

### Messages carried using the SOAP over JMS format:

A typical usage pattern scenario of the IBM MQ custom channel for WCF is as an interface to web services hosted over IBM MQ (SOAP/JMS).

Messages are carried using the SOAP over JMS message format of IBM MQ, enabling WCF clients and services to also call or be called by other WebSphere Application Server applications or hosting environments which are compatible with this format, including web services and clients running in , CICS, Axis v1 ( Java ), and .asmx (.NET), as shown in the following diagram:



For details on SOAP over JMS, see: "Developing web services with IBM MQ transport for SOAP" on page 1218

An example of a typical scenario from the diagram would be:

1. A web Service hosted within WebSphere Application Server and exposed over IBM MQ using the support for SOAP over JMS within WebSphere Application Server.
2. The WSDL document describing the service can then be used by the WCF tool to generate a client proxy and configuration which would then create an appropriate WCF channel stack, including the custom channel.
3. The client application can then use the proxy to start the web service in the same way as any other web service.

The channel would typically be used with a WCF text/SOAP message encoder, but the channel can be paired with other WCF message encoders if required. Using alternative encoders can also provide limited integration with native IBM MQ applications which do not support SOAP over JMS, but this is not the primary role of the channel.

The key benefits of using the custom channel in a WCF environment are:

- Asynchronous invocation: Supporting fire and forget client operations where the client is decoupled from the availability of the service and features, such as rerouting of responses and multi-hop.
- Reliable scaling characteristics: Queue based messaging allows capacity to be predictably added to a system.
- Quality of service: Messages are tangible and traceable, and can be easily managed and administered.

### **Messages carried using the Non-SOAP/Non-JMS message (Pure MQMessage) format:**

When you use the IBM MQ custom channel for WCF as a non-SOAP interface for the transmission of native IBM MQ messages, the messages are carried by using the Non-SOAP/Non-JMS message (Pure MQMessage) format of IBM MQ.

WCF users are able to start the service, or in other words, service users are able to send a message to an IBM MQ queue by using MQMessages. Applications can get and set the MQMD fields and payload. When the message is available in IBM MQ MQ queues, this message can be processed by any WCF service or non-WCF applications such as C or Java applications that are running on Windows, UNIX or z/OS.

### **Software requirements and installation instructions for the IBM MQ custom channel for WCF**

This topic outlines the software requirements and installation information for the IBM MQ custom channel for WCF.

The IBM MQ custom channel for WCF can only connect to IBM WebSphere MQ Version 7.0 or higher queue managers.

### **Software requirements for the WCF custom channel for IBM MQ**

This information lists the software requirements for the WCF custom channel for IBM MQ.

#### **Runtime environment**

- Microsoft.NET Framework v3.5 or higher must be installed on the host machine.
- *Java and .NET Messaging and Web Services* is installed by default as part of the IBM MQ Version 8.0 installer. Installs the .NET assemblies needed for the custom channel into the Global Assembly Cache.

**Note:** If the Microsoft.NET Framework v3.5 or higher is not installed before installing IBM MQ Version 8.0 or later, then the IBM MQ product installation continues without error, but the IBM MQ custom channel is unavailable. If the .NET Framework is installed after installing IBM MQ Version 8.0 or later, then the IBM MQ custom channel must be activated by running the `WMQInstallDir\bin\amqiRegisterdotNet.cmd` script, where `WMQInstallDir` is the directory where IBM MQ Version 8.0 or later is installed. This script installs the required assemblies in the Global Assembly Cache (GAC). A set of `amqi*.log` files recording the actions taken are created in the `%TEMP%` directory. It is not necessary to rerun the `amqiRegisterdotNet.cmd` script if .NET is upgraded to v3.5 or higher from an earlier version, for example, from .NET v2.0.

#### **Development environment**

- Microsoft Visual Studio 2008 or Windows Software Development Kit for .NET 3.5 or later.
- Microsoft.NET Framework V3.5 or higher must be installed on the host machine in order to build the sample solution files.

**Note:** If the Microsoft.NET Framework v3.5 or higher is not installed before installing IBM MQ Version 8.0 or later, then the IBM MQ product installation continues without error, but the IBM MQ custom channel is unavailable. If the .NET Framework is installed after installing IBM MQ Version 8.0 or later,

then the IBM MQ custom channel must be activated by running the `WMQInstallDir\bin\amqiRegisterdotNet.cmd` script, where `WMQInstallDir` is the directory where IBM MQ Version 8.0 or later is installed. This script installs the required assemblies in the Global Assembly Cache (GAC). A set of `amqi*.log` files recording the actions taken are created in the `%TEMP%` directory. It is not necessary to rerun the `amqiRegisterdotNet.cmd` script if .NET is upgraded to v3.5 or higher from an earlier version, for example, from .NET v2.0.

## IBM MQ custom channel for WCF: What's installed?

The custom channel for IBM MQ is a transport channel using the Microsoft Windows Communication Foundation (WCF) unified programming model. The custom channel is installed by default as part of the IBM MQ Version 8.0 or later installation.

## IBM MQ custom channel for WCF

The IBM MQ custom channel for WCF is installed by default as part of the IBM MQ Version 8.0 installation. The custom channel and its dependencies are contained within the Java and .NET Messaging and Web Services component, which is installed by default. When upgrading to IBM MQ Version 8.0 from an earlier version, the update installs the IBM MQ custom channel for WCF by default if the Java and .NET Messaging and Web Services component was previously installed in an earlier installation.

The .NET Messaging and Web Services component contains the `IBM.XMS.WCF.dll` file, and the `IBM.WMQ.WCF.dll` file, and these files are the main custom channel assembly, which contains the WCF interface classes. These files are installed in the Global Assembly Cache (GAC) and are also available in the following directory: `MQ_INSTALLATION_PATH\bin` where `MQ_INSTALLATION_PATH` is the directory in which IBM MQ is installed.

The following table summarizes the key classes that are required for using the custom channel.

Table 170. Key classes required for using the custom channel

	SOAP/JMS interface (Existing)	Non-SOAP/Non-JMS interface (From IBM MQ Version 8.0 )
Custom Channel Assembly	IBM.XMS.WCF.dll	IBM.WMQ.WCF.dll
Transport Binding Name	IBM.XMS.WCF.SoapJmsIbmTransportBindingElement	IBM.WMQ.WCF.WmqIbmTransportBindingElement
Transport Binding Importer	IBM.XMS.WCF.SoapJmsIbmTransportBindingElementImporter	IBM.WMQ.WCF.WmqIbmTransportBindingElementImporter
Transport Binding Config	IBM.XMS.WCF.SoapJmsIbmTransportBindingElementConfig	IBM.WMQ.WCF.WmqIbmTransportBindingElementConfig
Samples(Oneway)	SimpleOneWay_Client, SimpleOneWay_Service	MQMessaging_OneWay_Client, MQMessaging_OneWay_Service
Samples(RequestReply)	SimpleRequestReply_Client, SimpleRequestReply_Service	MQMessaging_RequestReply_Client, MQMessaging_RequestReply_Service

IBM.WMQ.WCF.dll supports both SOAP/JMS and Non-SOAP/Non-JMS interfaces. New applications developed are recommended to use the IBM.WMQ.WCF assembly as it supports both interfaces.

### ► V 9.0.5

## Sending MQSTR formatted messages

From IBM MQ Version 9.0.5, if the request message is of type MQSTR, you can select to send the reply message in MQSTR format.

You must use an additional URI parameter **replyMessageFormat** to change the format of the reply message. The supported values are:

"" " " is the default value.

The reply message is in byte (MQMFT\_NONE) format. For example:

```
"jms:/queue?destination=SampleQ@QM1&connectionFactory=binding(server)connectQueueManager(QM1)
&initialContextFactory=com.ibm.mq.jms.NoJndi&replyDestination=SampleReplyQ&replyMessageFormat= "
```

### MQSTR

The reply message is in MQSTR (MQMFT\_STRING) format. For example:

```
"jms:/queue?destination=SampleQ@QM1&connectionFactory=binding(server)connectQueueManager(QM1)
&initialContextFactory=com.ibm.mq.jms.NoJndi&replyDestination=SampleReplyQ&replyMessageFormat=MQSTR"
```

### Notes:

1. The value for **replyMessageFormat** is case insensitive.
2. Using any value other than " " or *MQSTR*, causes an invalid parameter value exception.

## IBM MQ custom channel samples

The samples provide some simple examples of how the IBM MQ custom channel for WCF can be used. The samples and their associated files are located in the *MQ\_INSTALLATION\_PATH* \tools\dotnet\samples\cs\wcf directory, where *MQ\_INSTALLATION\_PATH* is the installation directory for IBM MQ. For more information on the IBM MQ custom channel samples, see "Using the WCF samples" on page 1204.

### svcutil.exe.config

The *svcutil.exe.config* is an example of the configuration settings required to enable the Microsoft WCF *svcutil* client proxy generation tool to recognize the custom channel. The *svcutil.exe.config* file is located in the *MQ\_INSTALLATION\_PATH* \tools\wcf\docs\examples\ directory, where *MQ\_INSTALLATION\_PATH* is the installation directory for IBM MQ. For more information on using the *svcutil.exe.config*, see "Generating a WCF client proxy and application configuration files using the *svcutil* tool with metadata from a running service" on page 1201.

## WCF architecture

The IBM MQ custom channel for WCF is integrated on top of the IBM Message Service Client for .NET (XMS .NET) API.

## SOAP/JMS interface

The WCF architecture is as shown in the following diagram:

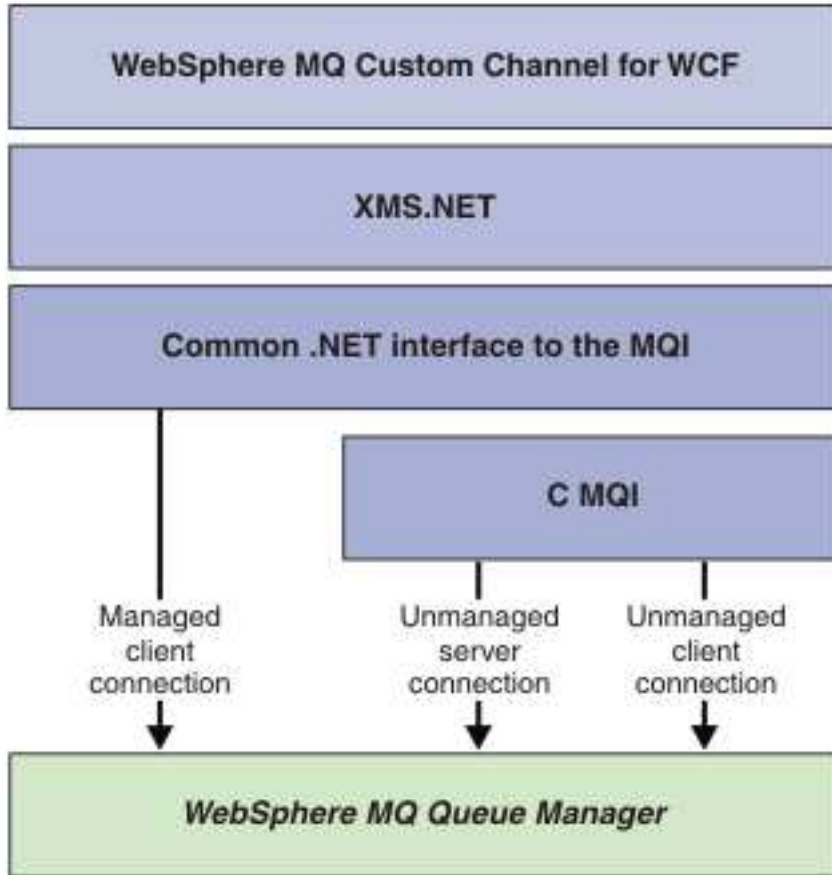


Figure 150. WCF architecture for the SOAP/JMS interface

For IBM WebSphere MQ Version 7.0.1 and later, all the required components are installed by default with the product installation.

The three connections are:

- Managed client connections
- Unmanaged server connections
- Unmanaged client connections

For more information about these connections, see “WCF Connection options” on page 1191.

### Non-SOAP/Non-JMS interface

The IBM MQ custom channel for WCF supports both the SOAP/JMS interface (available from IBM WebSphere MQ Version 7.0.1 ) and the Non-SOAP/Non-JMS interface.

The WCF architecture is as shown in the following diagram:

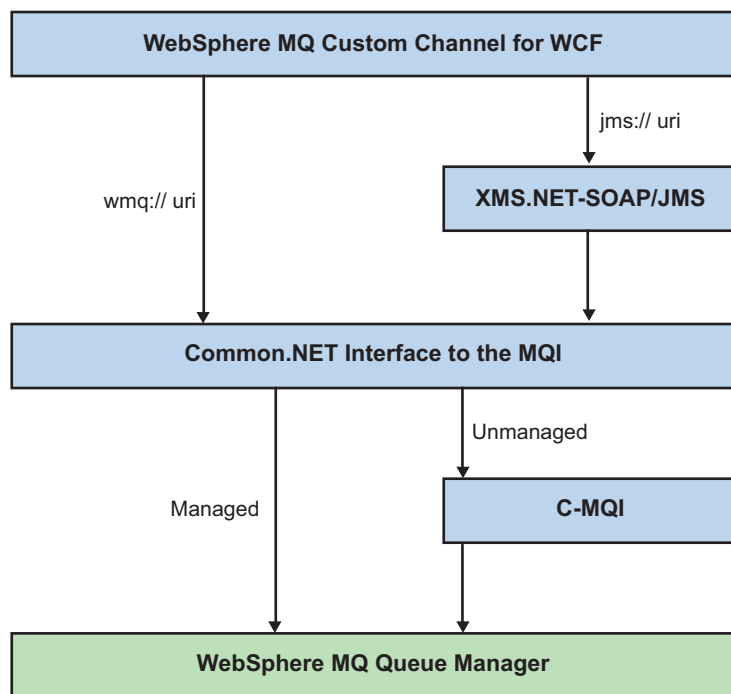


Figure 151. WCF architecture for the Non-SOAP/Non-JMS interface

## Using IBM MQ custom channels for WCF

Overview of the information available for programmers using IBM MQ custom channels for Windows Communication Foundation (WCF).

The Microsoft Windows Communication Foundation underpins the web services and messaging support in the Microsoft.NET Framework 3. IBM WebSphere MQ Version 7.0 or later can be used as a custom channel within WCF in the .NET Framework 3 in the same manner as the built-in channels offered by Microsoft.

Messages transported across the custom channel are formatted according to the SOAP over JMS implementation of IBM WebSphere MQ Version 7.0 or later. Applications can then communicate with services hosted by WCF or by the WebSphere SOAP over JMS service infrastructure. For more information about SOAP over JMS, see “Developing web services with IBM MQ transport for SOAP” on page 1218.



## WCF Custom channel features and capabilities

Use the following topics for information regarding WCF custom channel features and capabilities.

### WCF custom channel shapes:

Overview of the custom channel shapes that IBM MQ can be used as within the Microsoft Windows Communication Foundation (WCF) custom channels.

The IBM MQ custom channel for WCF supports two channel shapes:

- One-way
- Request-reply

WCF automatically selects the channel shape according to the service contract being hosted.

Contracts that include methods that only use the **IsOneWay** parameter are serviced by the one-way channel shape, for example:

```
[OperationContract(IsOneWay = true)]
void printString(String text);
```

Contracts that include either a mixture of one-way and request-reply methods, or all request-reply methods, are serviced by the request-reply channel shape. For example:

```
[OperationContract]
int subtract(int a, int b);

[OperationContract(IsOneWay = true)]
void printString(string text);
```

**Note:** When mixing one-way and request-reply methods in the same contract, you must ensure that the behavior is as intended, especially when working within a mixed environment because one-way methods wait until they receive a null reply from the service.

### One-way channel

The IBM MQ one-way custom channel for WCF is used, for example, to send messages from a WCF client using a one-way channel shape. The channel can send messages in one direction only, for example; from a client queue manager to a queue on a WCF service.

### Request-reply channel

The IBM MQ request-reply custom channel for WCF is used, for example, to send messages in two directions asynchronously; The same client instance must be used for asynchronous messaging. The channel can send messages in one direction, for example; from a client queue manager to a queue on a WCF service, and then send a reply message from the WCF to a queue on the client queue manager.

## WCF URI parameter names and values:

URI parameter names and values for the SOAP/JMS interface and Non-SOAP/Non JMS interface.

### SOAP/JMS interface

#### connectionFactory

The connectionFactory parameter is required. For the syntax of this parameter, see URI syntax and parameters for web service deployment.

#### initialContextFactory

The initialContextFactory parameter is required and must be set to "com.ibm.mq.jms.Nojndi" for compatibility with WebSphere Application Server and other products (see "Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP" on page 1275).

### Non-SOAP/Non JMS interface

The URI format is as for the MA93 specifications. See SupportPac - MA93 for further details of the IBM MQ IRI specifications.

#### IBM MQ URI syntax

```
wmq-iri = "wmq:" ["/" connection-name] "/" wmq-dest ["?" parm *("&" parm)]
connection-name = tcp-connection-name / other-connection-name
tcp-connection-name = ihost [":" port]
other-connection-name = 1*(iunreserved / pct-encoded)
wmq-dest = queue-dest / topic-dest
queue-dest = "msg/queue/" wmq-queue ["@" wmq-qmgr]
wmq-queue = wmq-name
wmq-qmgr = wmq-name
wmq-name = 1*48(wmq-char)
topic-dest = "msg/topic/" wmq-topic
wmq-topic = segment *("/" segment)
```

#### IBM MQ IRI example

The following example IRI tells a service requester that it can use a IBM MQ TCP client-binding connection to a machine called example.com on port 1414 and put persistent request messages to a queue called SampleQ on queue manager QM1. The IRI specifies that the service provider will put replies to a queue called SampleReplyQ.

```
1) wmq://example.com:1414/msg/queue/SampleQ@QM1?
ReplyTo=SampleReplyQ&persistence=MQPER_NOT_PERSISTENT
2) wmq://localhost:1414/msg/queue/Q1?
connectQueueManager=QM1&replyTo=Q2&connectionmode=managed
```

#### For TLS enabled connections

To make Secured (TLS) connections using the WCF Client/Service, set following properties with appropriate values in the URI. All the properties that are prefixed with "\*" are mandatory to make a secured connection.

- **sslKeyRepository:** \*SYSTEM or \*USER
- \* **sslCipherSpec:** a valid CipherSpec, for example TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA256.
- **sslCertRevocationCheck:** true or false.
- **sslKeyResetCount:** a value greater than 32kb.
- **sslPeerName:** the distinguished name of the server certificate

For example:

```
"wmq://localhost:1414/msg/queue/SampleQ?
connectQueueManager=QM1&sslkeyrepository=*SYSTEM&sslcipherspec=
TLS_RSA_WITH_AES_128_CBC_SHA&sslcertrevocationcheck=true&"sslpe
ername=" + " + "CN=ibmwebspheremqmm&sslkeyresetcount=45000"
```

## WCF custom channel assured delivery:


Assured Delivery guarantees that a service request or reply is actioned and not lost.

A request message is received and any reply message is sent under a local transaction sync point, which can be rolled back in the case of runtime failure. Examples of these failures are: An unhandled exception thrown by the service, failure to dispatch the message to the service, or failure to deliver the reply message.

AssuredDelivery is the assured delivery attribute which can be specified on a service contract to guarantee that any request messages received by a service, and any reply message sent from a service, is not lost in the event of a runtime failure.

To ensure that messages are also preserved in the event of system failure or power outage, messages must be sent as persistent. To use persistent messages the client application must have this option specified on its endpoint URI. For further information about setting URI properties, see: URI syntax and parameters for web service deployment.

Distributed transactions are not supported, and the scope of the transaction does not extend beyond the request and reply message processing performed by IBM MQ. Any work performed within the service might get rerun as a result of a failure which causes the message to be received again. The following diagram shows the scope of the transaction:



Assured delivery is enabled by applying the AssuredDelivery attribute to the service class as shown in the following example:

```
[AssuredDelivery]
class TestCalculatorService : IWMQSampleCalculatorContract
{
 public int add(int a, int b)
 {
 int ans = a + b;
 return ans;
 }
}
```

When using the `AssuredDelivery` attribute, you must be aware of the following points:

- When a channel determines that a failure is likely to recur if a message was rolled-back and received again, the message is treated as a poison message and is not returned to the request queue for reprocessing. For example: If the received message is not correctly formatted or cannot be dispatched to a service. Unhandled exceptions thrown from a service operation are always resent until the message has been redelivered the maximum number of times specified by the backout threshold property of the request queue. For more information, see: “WCF custom channel poison messages” on page 1189
- The channel performs the reading, processing, and replying of each request message as an atomic operation using a single thread of execution to enforce transactional integrity. To enable service operations to run concurrently, the channel enables WCF to create multiple instances of the channel. The number of channel instances available for processing requests is controlled by the binding property `MaxConcurrentCalls`. For more information, see: “WCF binding configuration options” on page 1197
- The assured delivery function uses both the `IOperationInvoker` and the `IErrorHandler` WCF extensibility points. If these extensibility points are used externally by an application, the application must ensure that any previously registered extensibility points are called. Failure to do so for `IErrorHandler` can result in errors going unreported. Failure to do so for `IOperationInvoker` can cause WCF to stop responding.

#### **WCF custom channel security:**

The IBM MQ custom channel for WCF supports the use of TLS only for unmanaged client connections to the queue manager.

TLS can be specified in one of two ways:

- Specify TLS directly on the SOAP over JMS URI. For a full description of the TLS options, see TLS and the IBM MQ transport for SOAP
- Specify TLS using an entry in the client channel definition table (CCDT). For more information about CCDTs, see Client channel definition table

#### **WCF client channel definition tables (CCDT):**

The IBM MQ custom channel for WCF supports the use of client channel definition tables (CCDT) to configure the connection information for client connections.

CCDTs are controlled through these two environment variables:

- `MQCHLLIB` specifies the directory where the table is located.
- `MQCHLTAB` specifies the file name of the table.

You cannot specify the channel definition table directly in the SOAP over JMS URI. If these environment variables are defined, then they take priority over any client connection details specified in the URI.

For more information about client channel definition tables, see: Client channel definition table.

**Related information:**

Client channel definition table

**WCF custom channel poison messages:**

When a service fails to process a request message, or fails to deliver a reply message to a reply queue, then the message is treated as a poison message.

**Poison request messages**

If a request message cannot be processed, then it is treated as a poison message. This action prevents the service from receiving the same unprocessable message again. For an unprocessable request message to be treated as a poison message, one of the following situations must be true:

- The messages backout count exceeded the backout threshold specified on the request queue, which only occurs if assured delivery was specified for the service. For more information about assured delivery, see: “WCF custom channel assured delivery” on page 1187
- The message was not formatted correctly and could not be interpreted as a SOAP over JMS message.

**Poison reply messages**

If a service fails to deliver a reply message to the reply queue, then the reply message is treated as a poison message. For reply messages, this action enables the reply messages to be retrieved later to aid problem determination.

**Poison message handling**

The action taken for a poison message depends on the queue manager configuration and the values set in the report options of the message. For SOAP over JMS, the following report options are set on request messages by default and are not configurable:

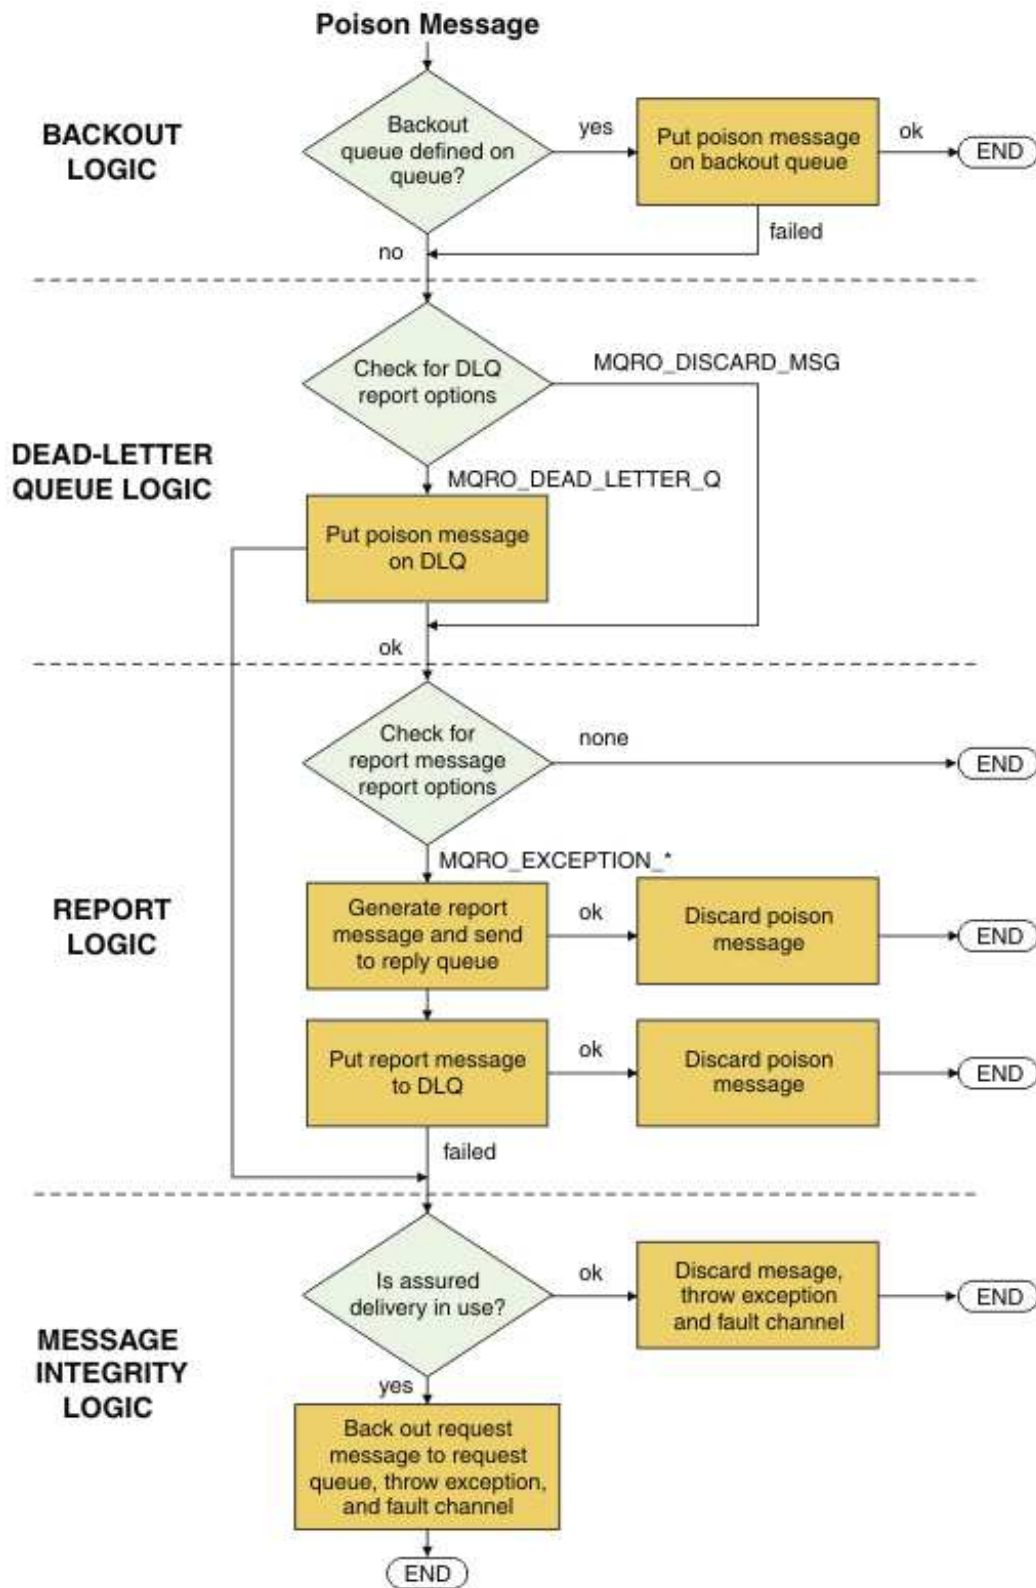
- MQRO\_EXCEPTION\_WITH\_FULL\_DATA
- MQRO\_EXPIRATION\_WITH\_FULL\_DATA
- MQRO\_DISCARD\_MSG

For SOAP over JMS, the following report option is set on reply messages by default and is not configurable:

- MQRO\_DEAD\_LETTER\_Q

If messages come from a non-WCF source, then refer to the documentation for that source.

The following diagram shows the possible actions and the steps taken if a poison message handling fails:



**IBM MQ message capabilities for WCF applications:**

Non-SOAP/Non-JMS (that is, IBM MQ ) message capabilities for WCF applications.

For the Non-SOAP/Non-JMS interface, the IBM MQ message capabilities for WCF applications are as follows:

- WCF applications can send and receive the base IBM MQ messages which can be processed by any IBM MQ application.
- WCF applications have full control to update the MQMD and payload.
- The WCF client can send IBM MQ messages that can be consumed by any IBM MQ clients, for example C, Java, JMS, and .NET clients.

The WCF for Non-SOAP/Non-JMS interface must use the following classes for setting the message payload and MQMD for the message:

- `WmqStringMessage` for a payload of type `String`
- `WmqBytesMessage` for a payload of type `Bytes`
- `WmqXmlMessage` for a payload of type `XML`

To set the payload of the message, use the **Data** property for the `WmqStringMessage`, `WmqBytesMessage` or `WmqXmlMessage` class, depending on the payload type. For example, use the following code to set a payload of type `String`:

```
WmqStringMessage strMsg = new WmqStringMessage();
//Setting the Message Payload
strMsg.Data = "Hello World";
//MQMD property
strMsg.Format = WmqMessageFormat.MQFMT_STRING;
```

## WCF Connection options

There are three modes of connecting an IBM MQ custom channel for WCF to a queue manager. Consider which type of connection best suits your requirements.

For more information about connection options, see: “Connection differences” on page 647

For more information about WCF architecture, see: “WCF architecture” on page 1182

## Unmanaged client connection

A connection made in this mode connects as an IBM MQ client to an IBM MQ server running either on the local machine or on a remote machine.

To use the IBM MQ custom channel for WCF as an IBM MQ client, you can install it, with the IBM MQ MQI client, either on the IBM MQ server, or on a separate machine.

## Unmanaged server connection

When used in server bindings mode, the IBM MQ custom channel for WCF uses the queue manager API, rather than communicating through a network. Using bindings connections provides better performance for IBM MQ applications than using network connections.

To use the bindings connection, you must install the IBM MQ custom channel for WCF on the IBM MQ server.

## Managed client connection

A connection made in this mode connects as an IBM MQ client to an IBM MQ server running either on the local machine or on a remote machine.

The IBM MQ custom channel classes for .NET 3 connecting in this mode remain in .NET managed code and make no calls to native services. For more information about managed code, see Microsoft documentation.

There are a number of limitations to using the managed client. For more information about these limitations, see “Managed client connections” on page 647.

## **Creating and configuring the IBM MQ custom channel for WCF**

The IBM MQ custom channels for WCF work in the same manner as transport WCF channels offered by Microsoft. The IBM MQ custom channel for WCF can be created in one of two ways.

### **About this task**

The IBM MQ custom channel integrates with WCF as a WCF transport channel, and as such must be paired with a message encoder and optional protocol channels, so it can create a complete channel stack that can be used by an application. Two elements are required for a complete channel stack to be created successfully:

1. A binding definition: Specifies which elements are required to build the applications channel stack, including transport channel, message encoder, and any protocols, plus any general configuration settings. For the custom channel, the binding definition must be created in the form of a WCF custom binding.
2. An endpoint definition: Links the service contract with the binding definition, and also provides the actual connection URI which describes where the application can connect. For the custom channel, the URI is in the form of a SOAP over JMS URI.

These definitions can be created in one of two different ways:

- Administratively; The definitions are created by providing the details in an application configuration file (for example: `app.config`).
- Programmatically; The definitions are created directly from the application code.

The decision over which method to use to create the definitions must be based on the requirements of the application as follows:

- The Administrative method for configuration provides the flexibility to alter the details of the service and client post-deployment without rebuilding the application.
- The Programmatic method for configuration provides greater protection from configuration errors, and the ability to dynamically generate a configuration at run time.

### **Creating a WCF custom channel administratively by supplying binding and endpoint information in an application configuration file:**

The IBM MQ custom channel for WCF is a transport level WCF channel. An endpoint and binding must be defined to use the custom channel, and these definitions can be done by supplying the binding and endpoint information in an application configuration file.

To configure and use the IBM MQ custom channel for WCF, which is a transport level WCF channel, a binding and an endpoint definition must be defined. The binding holds the configuration information for the channel, and the endpoint definition holds the connection details. These definitions can be created in two ways:

- Programmatically directly from the application code, as described here: “Creating a WCF custom channel by supplying binding and endpoint information programmatically” on page 1194
- Administratively, by providing the details in an application configuration file, as described in the following procedure.



The client or service application configuration file is commonly named *yourappname.exe.config* where *yourappname* is the name of your application. The application configuration file is most easily modified by using the Microsoft service configuration editor tool called *SvcConfigEditor.exe* in the following way:

- Start the *SvcConfigEditor.exe* configuration editor tool. The default installation location for the tool is: *Drive:\Program Files\Microsoft SDKs\Windows\v6.0\Bin\SvcConfigEditor.exe* where *Drive:* is the name of the installation drive.

### Step 1: Add a binding element extension to enable WCF to locate the custom channel

1. Right-click **Advanced > Extension > binding element** to open the menu, and select **New**
2. Complete the fields as shown in this table:

Table 171. New binding element fields

Field	Value
Name	IBM.XMS.WCF.SoopJmsIbmTransportChannel
Type	Navigate to IBM.XMS.WCF.dll in the Global Assembly Cache (GAC) and select IBM.XMS.WCFSoopJmsIbmTransportBindingElementConfig

### Step 2: Create a custom binding definition which pairs the custom channel with a WCF message encoder

1. Right-click **Bindings** to open the menu, and select **New Binding Configuration**
2. Complete the fields as shown in this table:

Table 172. New binding configuration fields

Field	Value
Name	CustomBinding_WMQ
BindingElement 1	textMessageEncoding (MessageVersion: Soap11)
BindingElement 2	IBM.XMS.WCF.SoopJmsIbmTransportChannel

### Step 3: Specify the binding properties

1. Select the *IBM.XMS.WCF.SoopJmsIbmTransportChannel* transport binding from the binding you created in: “Step 2: Create a custom binding definition which pairs the custom channel with a WCF message encoder”
2. Make any required changes to the default values of the properties as described in: “WCF binding configuration options” on page 1197

### Step 4: Create an endpoint definition

Create an endpoint definition which references the custom binding you created in: “Step 2: Create a custom binding definition which pairs the custom channel with a WCF message encoder” and provides the connection details of the service. The way this information is specified is dependent on whether the definition is for a client application or a service application.

For a client application, add an endpoint definition to the client section as follows:

1. Right-click **Client > Endpoints** to open the menu, and select **New Client Endpoint**
2. Complete the fields as shown in this table:

Table 173. New client endpoint fields

Field	Value
Name	Endpoint_WMQ
Address	The SOAP/JMS URI describing the WMQ connection details required to access the service. For further details see: "IBM MQ custom channel for WCF endpoint URI address format" on page 1196
Binding	customBinding
BindingConfiguration	CustomBinding_WMQ
Contract	The name of your service contract interface

For a service application, add a service definition to the services section as follows:

1. Right-click **Services** to open the menu, and select **New Service**, then select the service class to be hosted.
2. Add an endpoint definition to the Endpoints section for your new service, and complete the fields as shown in this table:

Table 174. New service endpoint fields

Field	Value
Name	Endpoint_WMQ
Address	The SOAP/JMS URI describing the WMQ connection details required to access the service. For further details see: "IBM MQ custom channel for WCF endpoint URI address format" on page 1196
Binding	customBinding
BindingConfiguration	CustomBinding_WMQ
Contract	The name of your service implementation class

### Creating a WCF custom channel by supplying binding and endpoint information programmatically:

The IBM MQ custom channel for WCF is a transport level WCF channel. An endpoint and binding must be defined to use the custom channel, and these definitions can be done programmatically directly from the application code.

To configure and use the IBM MQ custom channel for WCF, which is a transport level WCF channel, a binding and an endpoint definition must be defined. The binding holds the configuration information for the channel, and the endpoint definition holds the connection details. For more information see "Using the WCF samples" on page 1204.

These definitions can be created in two ways:

- Administratively, by providing the details in an application configuration file, as described in "Creating a WCF custom channel administratively by supplying binding and endpoint information in an application configuration file" on page 1192.
- Programmatically directly from the application code, as described in the following subtopics.

*Defining binding and endpoint information programmatically: SOAP/JMS interface:*

For the SOAP/JMS interface, you can define an endpoint and binding programmatically directly from the application code.

### **About this task**

To supply binding and endpoint information programmatically, add the required code to your application by completing the following steps.

### **Procedure**

1. Create an instance of the transport binding element of the channel by adding the following code to your application:  

```
SoapJmsIbmTransportBindingElement transportBindingElement = new SoapJmsIbmTransportBindingElement();
```
2. Set any required binding properties, for example, by adding the following code to your application to set the `ClientConnectionMode`:  

```
transportBindingElement.ClientConnectionMode = XmsWCFBindingProperty.AS_URI;
```
3. Create a custom binding that pairs the transport channel with a message encoder by adding the following code to your application:  

```
Binding binding = new CustomBinding(new TextMessageEncodingBindingElement(), transportBindingElement);
```
4. Create the SOAP/JMS URI. The SOAP/JMS URI that describes the IBM MQ connection details required to access the service, must be provided as the endpoint address. The address that you specify depends on whether the channel is being used for a service application or a client application.
  - For client applications, the SOAP/JMS URI must be created as an `EndpointAddress` as follows:  

```
EndpointAddress address = new EndpointAddress("jms:/queue?destination=SampleQ@QM1&connectionFactory=connectQueueManager(QM1)&initialContextFactory=com.ibm.mq.jms.Nojndi");
```
  - For service applications, the SOAP/JMS URI must be created as a `Uri` as follows:  

```
Uri address = new Uri("jms:/queue?destination=SampleQ@QM1&connectionFactory=connectQueueManager(QM1)&initialContextFactory=com.ibm.mq.jms.Nojndi");
```For more information about endpoint addresses, see “IBM MQ custom channel for WCF endpoint URI address format” on page 1196.

*Defining binding and endpoint information programmatically: Non-SOAP/Non-JMS interface:*

For the Non-SOAP/Non-JMS interface, you can define an endpoint and binding programmatically directly from the application code.

### **About this task**

To supply binding and endpoint information programmatically, add the required code to your application by completing the following steps.

### **Procedure**

1. Create a `WmqBinding` by adding the following code to your application:

```
WmqBinding binding = new WmqBinding();
```

This code creates a binding that pairs the `WmqMsgEncodingElement` and `WmqIbmTransportBindingElement` required for the Non-SOAP/Non-JMS interface.

2. Provide the `wmq://` URI that describes the IBM MQ connection details required to access the service. The way in which you provide the `wmq://` URI depends on whether the channel is being used for a service application or a client application.
  - For client applications, the `wmq://` URI must be created as an `EndpointAddress` as follows:

```
EndpointAddress address = new EndpointAddress
("wmq://localhost:1414/msg/queue/Q1?connectQueueManager=QM1&replyTo=Q2");
```

- For service applications, the wmq:// URI must be created as a URI as follows:

```
Uri sampleAddress = new Uri(
"wmq://localhost:1414/msg/queue/Q1?connectQueueManager=QM1&replyTo=Q2");
```

### IBM MQ custom channel for WCF endpoint URI address format:

A web service is specified using a Universal Resource Identifier (URI) that provides location and connection details. The URI format depends on whether you are using the SOAP/JMS interface or the Non-SOAP/Non-JMS interface.

#### SOAP/JMS interface

The URI format that is supported in the IBM MQ transport for SOAP permits a comprehensive degree of control over SOAP/ IBM MQ -specific parameters and options when accessing target services. This format is compatible with WebSphere Application Server and with CICS, facilitating the integration of IBM MQ with both those products.

The URI syntax is as follows:

```
jms:/queue? name=value&name=value...
```

where name is a parameter name and *value* is an appropriate value, and the name = *value* element can be repeated any number of times with the second and subsequent occurrences being preceded by an ampersand (&).

For further information about setting URI properties, see URI syntax and parameters for web service deployment

Parameter names are case-sensitive, as are names of IBM MQ objects. If any parameter is specified more than once, the final occurrence of the parameter takes effect meaning client applications can override parameter values by appending to the URI. If any additional unrecognized parameters are included, they are ignored.

If you store a URI in an XML string, you must represent the ampersand character as "&amp;". Similarly, if a URI is coded in a script, take care to escape characters such as & which would otherwise be interpreted by the shell.

This is an example of a simple URI for an Axis service:

```
jms:/queue?destination=myQ&connectionFactory=()
&initialContextFactory=com.ibm.mq.jms.NoJndi
```

Here is an example of a simple URI for a .NET service:

```
jms:/queue?destination=myQ&connectionFactory=()&targetService=MyService.asmx
&initialContextFactory=com.ibm.mq.jms.NoJndi
```

Only the required parameters are supplied ( targetService is required for .NET services only), and connectionFactory is given no options.

In this Axis example, connectionFactory contains a number of options:

```
jms:/queue?destination=myQ@myRQM&connectionFactory=connectQueueManager(myconnQM)
binding(client)clientChannel(myChannel)clientConnection(myConnection)
&initialContextFactory=com.ibm.mq.jms.NoJndi
```

In this Axis example, the sslPeerName option of connectionFactory has also been specified. The value of sslPeerName itself contains name value pairs and significant embedded blanks:

```

jms:/queue?destination=myQ@myRQM&connectionFactory=connectQueueManager(myconnQM)
binding(client)clientChannel(myChannel)clientConnection(myConnection)
sslPeerName(CN=MQ Test 1,O=IBM,S=Hampshire,C=GB)
&initialContextFactory=com.ibm.mq.jms.Nojndi

```

## NON-SOAP/Non-JMS interface

The URI format for the NON-SOAP/Non-JMS interface permits a comprehensive degree of control over IBM MQ -specific parameters and options when accessing target services.

The URI syntax is as follows:

```
wmq://example.com:1415/msg/queue/INS.QUOTE.REQUEST@MOTOR.INS ?ReplyTo=msg/queue/INS.QUOTE.REPLY@BRANCH452&persistence=MQPER_NOT_F
```

This IRI tells a service requester that it can use an IBM MQ TCP client-binding connection to a machine called example.com on port 1415 and put persistent request messages to a queue called INS.QUOTE.REQUEST on queue manager MOTOR.INS. The IRI specifies that the service provider puts replies to a queue called INS.QUOTE.REPLY on queue manager BRANCH452. The URI format is as specified for SupportPac MA93. See SupportPac MA93: IBM MQ - Service Definition for more details about the IBM MQ IRI specifications.

## WCF binding configuration options:

There are two ways of applying configuration options to the custom channels binding information. You either set the properties administratively, or set them programmatically.

The binding configuration options can be set in one of two different ways:

1. Administratively: The binding property settings must be specified in the transport section of the custom binding definition in the applications configuration file, for example: app.config.
2. Programmatically: The application code must be modified to specify the property during initialization of the custom binding.

## Setting the binding properties administratively

The binding property settings can be specified in the application configuration file, for example: app.config. The configuration file is generated by **svcutil**, as shown in the following examples.

### SOAP/JMS interface

```

<customBinding>
...
 <IBM.XMS.WCF.SoapJmsIbmTransportChannel maxBufferSize="524288"
 maxMessageSize="4000000" clientConnectionMode="0" maxConcurrentCalls="16"/>
...
</customBinding>

```

### Non-SOAP/Non-JMS interface

```

<customBinding>
 <IBM.WMQ.WCF.WmqMsgEncodingElement/>
 <IBM.WMQ.WCF.WmqIbmTransportChannel maxBufferSize="524288"
 maxMessageSize="65536" clientConnectionMode="managedclient"/>
</customBinding>

```

## Setting the binding properties programmatically

To add a WCF binding property to specify the client connection mode, you must modify the service code to specify the property during initialization of the custom binding.

Use the following example to specify unmanaged client connection mode:

```

SoapJmsIbmTransportBindingElement
transportBindingElement = new SoapJmsIbmTransportBindingElement();
transportBindingElement.ClientConnectionMode = XmsWCFBindingProperty.CLIENT_UNMANAGED;

Binding sampleBinding = new CustomBinding(new TextMessageEncodingBindingElement(),
 transportBindingElement);

```

## WCF binding properties

Table 175. Values of binding properties when setting administratively or programmatically

| Property name        | Client or Service application | Administrative value                                      | Programmatic value                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------|-------------------------------|-----------------------------------------------------------|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| maxBufferPoolSize    | Both                          | 0 to 64 bit signed integer                                | 0 to 64 bit signed integer                                | Specifies the maximum size of the memory that can be used to store WCF message buffers for an instance of the channel.                                                                                                                                                                                                                                                                                                                                    |
| maxMessageSize       | Both                          | 1 to 32 bit signed integer                                | 1 to 32 bit signed integer                                | Specifies the maximum memory that can be used for an individual WCF message.                                                                                                                                                                                                                                                                                                                                                                              |
| clientConnectionMode | Both                          | 0 (Default value)<br>1                                    | AS_URI (Default value)<br>CLIENT_UNMANAGED                | Specifies the client connection mode of the transport channel.<br><br>0 means that the client connection mode is as specified in the URI. Only used if the client connection is used. Specifies that the client connection mode is as specified in the URI. 0 is the default value if no client connection mode is set.<br><br>1 means that the client connection mode is an unmanaged client. Only used if the client connection is used.                |
| MaxConcurrentCalls   | Client                        | The range is 0 - 2 147 483 647<br>16 is the default value | The range is 0 - 2 147 483 647<br>16 is the default value | This property defines the maximum number of concurrent operations that can take place on an individual client proxy at any one time. If more operations are started, they are queued until an in-progress operation either completes or times out. This setting can be used to control the maximum threads and resources which can be consumed by an individual proxy.<br><br>0 removes this limit, enabling all operations to be attempted concurrently. |

Table 175. Values of binding properties when setting administratively or programmatically (continued)

| Property name      | Client or Service application | Administrative value                                          | Programmatic value                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|-------------------------------|---------------------------------------------------------------|---------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MaxConcurrentCalls | Service                       | The range is 1 - 2 147 483 647<br><br>16 is the default value | The range is 1 - 2 147 483 647<br><br>16 is the default value | This property is only used if the assured delivery feature is enabled (For more information about assured delivery, see “WCF custom channel assured delivery” on page 1187 ). It specifies the maximum number of concurrent operations that can be in progress at the same time for the given endpoint.<br><br>Care is needed when changing this setting. Each concurrent operation requires additional resources, in particular a new instance of the custom channel and the associated threads from the thread pool to action the requests. Over-allocating can be counter productive and affect performance severely. Appropriate configuration of the thread pool must be made to support this property. |

## Building and hosting services for WCF

Overview of Microsoft Windows Communication Foundation (WCF) services explaining how to create and configure WCF services.

The IBM MQ custom channel for WCF and the WCF services which use it, can be hosted by the following methods:

- Self-hosting
- Windows Service

The IBM MQ custom channel for WCF cannot be hosted in Windows Process Activation Service.

The following topics provide some simple self-hosting examples to demonstrate the steps involved. The Microsoft WCF online documentation, which contains further information and the latest details, can be found on the Microsoft MSDN website at <http://msdn.microsoft.com>.

## Building WCF service applications using method 1: Self-hosting administratively using an application configuration file:

Having created an application configuration file, open an instance of the service and add the specified code to your application.

### Before you begin

Create or edit an application configuration file for the service, as described in: "Creating a WCF custom channel administratively by supplying binding and endpoint information in an application configuration file" on page 1192

### About this task

1. Instantiate and open an instance of the service in the service host. The service type must be the same as the service type specified in the service configuration file.
2. Add the following code to your application:

```
ServiceHost service = new ServiceHost(typeof(MyService));
service.Open();
...
service.Close();
```

## Building WCF service applications using method 2: Self-hosting programmatically directly from the application:

Add the binding properties, create the service host with an instance of the required service class and open the service.

### Before you begin

1. Add a reference to the custom channel `IBM.XMS.WCF.dll` file to the project. The `IBM.XMS.WCF.dll` is in the `WMQInstallDir\bin` where `WMQInstallDir` is the directory that IBM MQ is installed in.
2. Add a `using` statement to the `IBM.XMS.WCF` namespace, for example: `using IBM.XMS.WCF`
3. Create an instance of the channels binding element and endpoint as described in: "Creating a WCF custom channel by supplying binding and endpoint information programmatically" on page 1194

### About this task

If changes to the binding properties of the channel are required, then complete the following steps:

1. Add the binding properties to `transportBindingElement` as shown in the following example:

```
SoapJmsIbmTransportBindingElement transportBindingElement = new SoapJmsIbmTransportBindingElement();
Binding binding = new CustomBinding(new TextMessageEncodingBindingElement(), transportBindingElement);
Uri address = new Uri("jms:/queue?destination=SampleQ@QM1&connectionFactory=connectQueueManager(QM1)&initialContextFactory=com.ibm.mq.jms.Nojndi");
```
2. Create the service host with an instance of the required service class:

```
ServiceHost service = new ServiceHost(typeof(MyService));
```
3. Open the service:

```
service.AddServiceEndpoint(typeof(IMyServiceContract), binding, address);
service.Open();
...
service.Close();
```



## Exposing metadata using an HTTP endpoint:

Instructions for exposing the metadata of a service which is configured to use the IBM MQ custom channel for WCF.

### About this task

If the services metadata must be exposed (so that tools such as `svcutil` can access it directly from the running service rather than from an offline WSDL file for example) it must be done by exposing the services metadata with an HTTP endpoint. The following steps can be used to add this additional endpoint.

1. Add the base address of where the metadata must be exposed to the `ServiceHost`, for example:

```
ServiceHost service = new ServiceHost(typeof(TestService),
 new Uri("http://localhost:8000/MyService"));
```

2. Add the following code to the `ServiceHost` before the service is opened:

```
ServiceMetadataBehavior metadataBehavior = new ServiceMetadataBehavior();
metadataBehavior.HttpGetEnabled = true;
service.Description.Behaviors.Add(metadataBehavior);
service.AddServiceEndpoint(typeof(IMetadataExchange),
 MetadataExchangeBindings.CreateMexHttpBinding(), "mex");
```

### Results

The metadata is now available at the following address: `http://localhost:8000/MyService`

## Building client applications for WCF

Overview of generating and building Microsoft Windows Communication Foundation (WCF) client applications.

A client application can be created for a WCF service; client applications are typically generated by using the Microsoft ServiceModel Metadata Utility Tool (`Svcutil.exe`) to create the required configuration and proxy files which can be used directly by the application.

### Generating a WCF client proxy and application configuration files using the `svcutil` tool with metadata from a running service:

Instructions for using the Microsoft `svcutil.exe` tool to generate a client for a service which is configured to use the IBM MQ custom channel for WCF.

### Before you begin

There are three prerequisites for using the `svcutil` tool to create to required configuration and proxy files which can be used directly by the application:

- The WCF service must be running before the `svcutil` tool is started.
- The WCF service must expose its metadata using an HTTP port in addition to the IBM MQ custom channel endpoint references to generate a client directly from a running service.
- The custom channel must be registered in the configuration data for `svcutil`.

### About this task

The following steps explain how to generate a client for a service which is configured to use the IBM MQ custom channel, but also exposes its metadata at run time via a separate HTTP port:

1. Start the WCF service (The service must be running before the `svcutil` tool is started).

2. Add the details from the `svcutil.exe` config file from the root of the installation, into the active `svcutil` configuration file, typically `C:\Program Files\Microsoft SDKs\Windows\v6.0A\bin\svcutil.exe.config` so `svcutil` recognizes the IBM MQ custom channel.
3. Run `svcutil` from a command prompt, for example:
 

```
svcutil /language:C# /r: installlocation\bin\IBM.XMS.WCF.dll
/config:app.config http://localhost:8000/IBM.XMS.WCF/samples
```
4. Copy the generated `app.config` and `YourService.cs` files to the Microsoft Visual studio client project.

## What to do next

If the services metadata cannot be directly retrieved, `svcutil` can be used to generate the client files from `wSDL` instead. For more information see: "Generating a WCF client proxy and application configuration files using the `svcutil` tool with WSDL"

## Generating a WCF client proxy and application configuration files using the `svcutil` tool with WSDL:

Instructions for generating WCF clients from WSDL if the metadata of the service is unavailable.

If the metadata of the service cannot be directly retrieved to generate a client from the metadata from a running service, then, `svcutil` can be used to generate the client files from WSDL instead. The following modifications must be made to the WSDL to specify that the IBM MQ custom channel is to be used:

1. Add the following namespace definitions and policy information:

```
<wsdl:definitions
xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy"
xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-utility-1.0.xsd">

 <wsp:Policy wsu:Id="CustomBinding_IWMQSampleContract_policy">
 <wsp:ExactlyOne>
 <wsp>All>
 <xms:xms xmlns:xms="http://sample.schemas.ibm.com/policy/xms" />
 </wsp>All>
 </wsp:ExactlyOne>
 </wsp:Policy>

 ...

</wsdl:definitions>
```

2. Modify the bindings section to refer to the new policy section and remove any transport definition from the underlying binding element:

```
<wsdl:definitions ...>

 <wsdl:binding ...>
 <wsp:PolicyReference URI="#CustomerBinding_IWMQSampleContract_policy" />
 <[soap]:binding ... transport="" />
 ...
 </wsdl:binding>
</wsdl:definitions>
```

3. Run `svcutil` from a command prompt, for example:

```
svcutil /language:C# /r: MQ_INSTALLATION_PATH\bin\IBM.XMS.WCF.dll
/config:app.config MQ_INSTALLATION_PATH\src\samples\WMQAxis\default\service
\soap.server.stockQuoteAxis_wmq.wsdl
```

## Building WCF client applications using a client proxy with an application configuration file:

### Before you begin

Create or edit an application configuration file for the client, as described in: “Creating a WCF custom channel administratively by supplying binding and endpoint information in an application configuration file” on page 1192

### About this task

Instantiate and open an instance of the client proxy. The parameter passed to the generated proxy must be the same as the endpoint name specified in the client configuration file, for example `Endpoint_WMQ`:

```
MyClientProxy myClient = new MyClientProxy("Endpoint_WMQ");
 try {
 myClient.myMethod("HelloWorld!");
 myClient.Close();
 }
 catch (TimeoutException e) {
 Console.Out.WriteLine(e);
 myClient.Abort();
 }
 catch (CommunicationException e) {
 Console.Out.WriteLine(e);
 myClient.Abort();
 }
 catch (Exception e) {
 Console.Out.WriteLine(e);
 myClient.Abort();
 }
}
```

## Building WCF client applications using a client proxy with programmatic configuration:

### Before you begin

1. Add a reference to the custom channel `IBM.XMS.WCF.dll` file to the project. The `IBM.XMS.WCF.dll` is in the `WMQInstallDir\bin` directory where `WMQInstallDir` is the directory that IBM MQ is installed in.
2. Add a `using` statement to the `IBM.XMS.WCF` namespace, for example: `using IBM.XMS.WCF`
3. Create an instance of the binding element and endpoint of the channel as described in: “Creating a WCF custom channel by supplying binding and endpoint information programmatically” on page 1194

### About this task

If changes to the binding properties of the channel are required, complete the following steps.

1. Add the binding properties to `transportBindingElement` as shown in the following figure:

```
SoapJmsIbmTransportBindingElement transportBindingElement = new SoapJmsIbmTransportBindingElement();
Binding binding = new CustomBinding(new TextMessageEncodingBindingElement(), transportBindingElement);
EndpointAddress address =
 new EndpointAddress("jms:/queue?destination=SampleQ@QM1&connectionFactory=
 connectQueueManager(QM1)&initialContextFactory=com.ibm.mq.jms.Nojndi");
```

2. Create the client proxy as shown in the following figure, where *binding* and *endpoint address* are the binding and endpoint address configured in step 1 and passed in:

```
MyClientProxy myClient = new MyClientProxy(binding, endpoint address);
 try {
 myClient.myMethod("HelloWorld!");
 myClient.Close();
 }
 catch (TimeoutException e) {
 Console.Out.WriteLine(e);
 myClient.Abort();
 }
}
```

```

 catch (CommunicationException e) {
 Console.Out.WriteLine(e);
myClient.Abort();
 }
 catch (Exception e) {
 Console.Out.WriteLine(e);
 myClient.Abort();
 }
 }
}

```

## Using the WCF samples

The Windows Communication Foundation (WCF) samples provide some simple examples of how the IBM MQ custom channel can be used.

To build the sample projects, either the Microsoft.NET 3.5 SDK, or Microsoft Visual Studio 2008 is needed.

### Simple one-way client and server WCF sample

This sample demonstrates the IBM MQ custom channel being used to start a Windows Communication foundation (WCF) service from a WCF client using a one-way channel shape.

#### About this task

The service implements a single method which outputs a string to the console. The client has been generated by using the `svcutil` tool to retrieve the service metadata from a separately exposed HTTP endpoint as described in “Generating a WCF client proxy and application configuration files using the `svcutil` tool with metadata from a running service” on page 1201

The sample has been configured with specific resource names as described in the following procedure. If you must change the resource names, then you must also change the corresponding value on the client application in the `MQ_INSTALLATION_PATH\tools\dotnet\samples\cs\wcf\samples\WCF\oneway\client\app.config` file, and on the service application in the `MQ_INSTALLATION_PATH\tools\dotnet\samples\cs\wcf\samples\WCF\oneway\service\TestServices.cs` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM MQ. For more information about formatting the JMS endpoint URI, see *IBM MQ Transport for SOAP* in the IBM MQ product documentation. If you need to modify the sample solution and source, then you need an IDE, for example, Microsoft Visual Studio 8 or higher.

#### Procedure

1. Create a queue manager called *QM1*
2. Create a queue destination called *SampleQ*
3. Start the service so the listener is waiting for messages: Run the `MQ_INSTALLATION_PATH\tools\dotnet\samples\cs\wcf\samples\WCF\oneway\service\bin\Release\TestService.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM MQ.
4. Run the client once: Run the `MQ_INSTALLATION_PATH\tools\dotnet\samples\cs\wcf\samples\WCF\oneway\client\bin\Release\TestClient.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM MQ. The client application loops five times sending five messages to *SampleQ*

#### Results

The service application gets the messages from *SampleQ* and displays Hello World on the screen five times.

#### What to do next

## Simple request-reply client and server WCF sample

This sample demonstrates the IBM MQ custom channel being used to start a Windows Communication foundation (WCF) service from a WCF client using a request-reply channel shape.

### About this task

This service provides some simple calculator methods to add and subtract two numbers, and then return the result. The client has been generated by using the `svcutil` tool to retrieve the service metadata from a separately exposed HTTP endpoint as described in “Generating a WCF client proxy and application configuration files using the `svcutil` tool with metadata from a running service” on page 1201

The sample has been configured with specific resource names as in the following procedure described. If you need to change the resource names, then you also need to change the corresponding value on the client application in the `MQ_INSTALLATION_PATH\Tools\wcf\samples\WCF\requestreply\client\app.config` file, and on the service application in the `MQ_INSTALLATION_PATH\Tools\wcf\samples\WCF\requestreply\service\RequestReplyService.cs` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM MQ. For more information about formatting the JMS endpoint URI, see *IBM MQ Transport for SOAP* in the IBM MQ product documentation. If you need to modify the sample solution and source, then you need an IDE, for example, Microsoft Visual Studio 8 or higher.

### Procedure

1. Create a queue manager called *QM1*
2. Create a queue destination called *SampleQ*
3. Create a queue destination called *SampleReplyQ*
4. Start the service so the listener is waiting for messages: Run the `MQ_INSTALLATION_PATH\Tools\wcf\samples\WCF\requestreply\service\bin\Release\SimpleRequestReply_Service.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM MQ.
5. Run the client once: Run the `MQ_INSTALLATION_PATH\Tools\wcf\samples\WCF\requestreply\client\bin\Release\SimpleRequestReply_Client.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM MQ.

### Results

When the client has been run, the following process is started and repeats four times so a total of five messages are sent each way:

1. The client puts a request message on *SampleQ* and waits for a response.
2. The service gets the request message from *SampleQ*.
3. The service adds and subtracts some values using the contents of the message.
4. The service then puts the results into a message on *SampleReplyQ*, and waits for the client to put a new message.
5. The client gets the message from *SampleReplyQ* and displays the results on the screen.

### What to do next

## WCF client to a .NET service hosted by IBM MQ sample

Sample client applications and sample service proxy applications are supplied for both .NET and Java. The samples are based on a Stock Quote service that takes a request for a stock quote and then provides the stock quote.

### Before you begin

The sample requires that the .NET SOAP over JMS service hosting environment is correctly installed and configured in IBM MQ and is accessible from a local queue manager. For information about installing and configuring the environment, see: “Installing IBM MQ Web transport for SOAP” on page 1229

When the .NET SOAP over JMS service hosting environment is correctly installed and configured in IBM MQ and is accessible from a local queue manager, additional configuration steps must be completed.

1. Set the WMQSOAP\_HOME environment variable to the IBM MQ installation directory, for example:  
C:\Program Files\IBM\MQ
2. Ensure that the Java compiler javac is available and on the PATH.
3. Copy the file axis.jar from the prereqs/axis directory of the WebSphere installation CD to the IBM MQ production directory, for example: C:\Program Files\IBM\MQ\java\lib\soap
4. Add to the PATH: *MQ\_INSTALLATION\_PATH*\Java\lib where *MQ\_INSTALLATION\_PATH* represents the directory where IBM MQ is installed, for example: C:\Program Files\IBM\MQ
5. Ensure that the location of .NET is specified correctly in *MQ\_INSTALLATION\_PATH*\bin\amqwallWSDL.cmd where *MQ\_INSTALLATION\_PATH* represents the directory where IBM MQ is installed, for example: C:\Program Files\IBM\MQ. The location of .NET can be specified for example: set msfwdir=%ProgramFiles%\Microsoft Visual Studio .NET 2003\SDK\v1.1\Bin

When the previous steps are complete, test and run the service:

1. Navigate to your SOAP over JMS working directory.
2. Enter one of the following commands to run the verification test and leave the service listener running:
  - For .NET: *MQ\_INSTALLATION\_PATH*\Tools\soap\samples\runivt dotnet hold where *MQ\_INSTALLATION\_PATH* represents the directory where IBM MQ is installed.
  - For AXIS: *MQ\_INSTALLATION\_PATH*\Tools\soap\samples\runivt Dotnet2AxisClient hold where *MQ\_INSTALLATION\_PATH* represents the directory where IBM MQ is installed.

The hold argument keeps the listeners running after the test completes.

If errors are reported during this configuration, you can remove all the changes so that the procedure may be restarted in the following way:

1. Delete the generated SOAP over JMS directory.
2. Delete the queue manager.

### About this task

This sample demonstrates a connection from a WCF client to the .NET SOAP over JMS sample service provided in IBM MQ using a one-way channel shape. The service implements a simple StockQuote example, which outputs a text string to the console.

The client has been generated by using WSDL to generate client files as described in “Generating a WCF client proxy and application configuration files using the svcutil tool with WSDL” on page 1202

The sample has been configured with specific resource names as described in the following procedure. If you need to change the resource names, then you must also change the corresponding value on the client application in the *MQ\_INSTALLATION\_PATH* \tools\wcf\samples\WMQNET\default\client\app.config file,

and on the service application in the `MQ_INSTALLATION_PATH\tools\wcf\samples\WMQNET\default\service\WmqDefaultSample_StockQuoteDotNet.wsd1` file, where `MQ_INSTALLATION_PATH` represents the installation directory for IBM MQ. For more information about formatting the JMS endpoint URI, see *IBM MQ Transport for SOAP* in the IBM MQ product documentation.

## Procedure

Run the client once: Run the `MQ_INSTALLATION_PATH\tools\wcf\samples\WMQNET\default\client\bin\Release\TestClient.exe` file, where `MQ_INSTALLATION_PATH` represents the installation directory for IBM MQ. The client application loops five times sending five messages to the sample queue.

## Results

The service application gets the messages from the sample queue and displays Hello World five times on the screen.

## WCF client to an Axis Java service hosted by IBM MQ sample

Sample client applications and sample service proxy applications are supplied for both Java and .NET. The samples are based on a Stock Quote service that takes a request for a stock quote and then provides the stock quote.

## Before you begin

This sample requires that the .NET SOAP over JMS service hosting environment is correctly installed and configured in IBM MQ and is accessible from a local queue manager. For information about installing and configuring the environment, see: “Installing IBM MQ Web transport for SOAP” on page 1229

When the .NET SOAP over JMS service hosting environment is correctly installed and configured in IBM MQ and is accessible from a local queue manager, additional configuration steps must be completed.

1. Set the `WMQSOAP_HOME` environment variable to the IBM MQ installation directory, for example:  
`C:\Program Files\IBM\MQ`
2. Ensure that the Java compiler `javac` is available and on the `PATH`.
3. Copy the file `axis.jar` from the `prereqs/axis` directory of the WebSphere installation CD to the IBM MQ installation directory.
4. Add to the `PATH`: `MQ_INSTALLATION_PATH\Java\lib` where `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed, for example: `C:\Program Files\IBM\MQ`
5. Ensure that the location of .NET is specified correctly in `MQ_INSTALLATION_PATH\bin\amqwca11WSDL.cmd` where `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed, for example: `C:\Program Files\IBM\MQ`. The location of .NET can be specified for example: `set msfwdir=%ProgramFiles%\Microsoft Visual Studio .NET 2003\SDK\v1.1\Bin`

When the previous steps are complete, test and run the service:

1. Navigate to your SOAP over JMS working directory.
2. Enter one of the following commands to run the verification test and leave the service listener running:
  - For .NET: `MQ_INSTALLATION_PATH\Tools\soap\samples\runivt dotnet hold` where `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed.
  - For AXIS: `MQ_INSTALLATION_PATH\Tools\soap\samples\runivt Dotnet2AxisClient hold` where `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed.

The `hold` argument keeps the listeners running after the test completes.

If errors are reported during this configuration, you can remove all the changes so that the procedure is restarted in the following way:

1. Delete the generated SOAP over JMS directory.
2. Delete the queue manager.

### About this task

The sample demonstrates a connection from a WCF client to the Axis Java SOAP over JMS sample service provided in IBM MQ using a one-way channel shape. The service implements a simple StockQuote example, which outputs a text string to a file which is saved in the current directory.

The client has been generated by using WSDL to generate client files as described in “Generating a WCF client proxy and application configuration files using the svcutil tool with WSDL” on page 1202

The sample has been configured with specific resource names as described in this paragraph. If you need to change the resource names, then you must also change the corresponding value on the client application in the `MQ_INSTALLATION_PATH\tools\wcf\samples\WMQAxis\default\client\app.config` file, and on the service application in the `MQ_INSTALLATION_PATH\tools\wcf\samples\WMQAxis\default\service\WmqDefaultSample_StockQuoteDotNet.wsdl` file, where `MQ_INSTALLATION_PATH` represents the installation directory for IBM MQ.

### Procedure

Run the client once: Run the `MQ_INSTALLATION_PATH\tools\wcf\samples\WMQAxis\default\client\bin\Release\TestClient.exe` file, where `MQ_INSTALLATION_PATH` represents the installation directory for IBM MQ. The client application loops five times sending five messages to the sample queue.

### Results

The service application gets the messages from the sample queue and adds Hello World five times to a file in the current directory.

#### Related reference:

“Handling different SOAP response element names” on page 1217

WCF expects the name of a returned value to be in a specific format by default, but a service might not return an element with its name in the expected format.

### WCF client to Java service hosted by WebSphere Application Server sample

Sample client applications and sample service proxy applications are supplied for WebSphere Application Server Version 6. A request-response service is also provided.

### Before you begin

This sample requires that the following IBM MQ configuration is used:

Table 176. IBM MQ required configuration

| Object        | Required name   |
|---------------|-----------------|
| Queue manager | QM1             |
| Local queue   | HelloWorld      |
| Local queue   | HelloWorldReply |

This sample also requires that a WebSphere Application Server Version 6 hosting environment is correctly installed and configured. WebSphere Application Server Version 6 uses a bindings mode connection to connect to IBM MQ by default. Therefore WebSphere Application Server Version 6 must be installed on the same machine as the queue manager.



After the WAS environment is configured, the following additional configuration steps must be completed:

1. Create the following JNDI objects in the WebSphere Application Server JNDI repository:
  - a. A JMS queue destination called HelloWorld
    - Set the JNDI name to `jms/HelloWorld`
    - Set the queue name to `HelloWorld`
  - b. A JMS queue connection factory called HelloWorldQCF
    - Set the JNDI name to `jms/HelloWorldQCF`
    - Set the queue manager name to `QM1`
  - c. A JMS queue connection factory called WebServicesReplyQCF
    - Set the JNDI name to `jms/WebServicesReplyQCF`
    - Set the queue manager name to `QM1`
2. Create a Message Listener Port called HelloWorldPort in WebSphere Application Server with the following configuration:
  - Set the connection factory JNDI name to `jms/HelloWorldQCF`
  - Set the destination JNDI name to `jms/HelloWorld`
3. Install the web service HelloWorldEJBEAR.ear application to your WebSphere Application Server as follows:
  - a. Click **Applications > New Application > New Enterprise Application**.
  - b. Navigate to `MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\HelloWorldsEJBEAR.ear` where `MQ_INSTALLATION_PATH` is the installation directory of IBM MQ.
  - c. Do not change any of the default option in the wizard, and restart the application server after the application has been installed.

When the WAS configuration is complete, test the service by running it once:

1. Navigate to your Soap over JMS working directory.
2. Enter this command to run the sample: `MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\TestClient.exe` where `MQ_INSTALLATION_PATH` is the installation directory of IBM MQ.

## About this task

The sample demonstrates a connection from a WCF client to the WebSphere Application Server SOAP over JMS sample service provided in the WCF samples included in IBM MQ, using a request-response channel shape. Messages flow between WCF and the WebSphere Application Server using IBM MQ queues. The service implements the `HelloWorld(...)` method, which takes a string and returns a greeting to the client.

The client has been generated by using the `svcutil` tool to retrieve the service metadata from a separately exposed HTTP endpoint as described in “Generating a WCF client proxy and application configuration files using the `svcutil` tool with metadata from a running service” on page 1201

The sample has been configured with specific resource names as described in the following procedure. If you need to change the resource names, then you must also change the corresponding value on the client application in the `MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\default\client\app.config` file, and on the service application in the `MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\HelloWorldsEJBEAR.ear` where `MQ_INSTALLATION_PATH` is the installation directory of IBM MQ. For more information about formatting the JMS endpoint URI, see URI syntax and parameters for web service deployment.

The service and client are based upon the service and client outlined in the developerWorks® article *Building a JMS web service using SOAP over JMS and WebSphere Studio*. For more information about developing SOAP over JMS web services that are compatible with the IBM MQ WCF custom channel, see

[http://www.ibm.com/developerworks/websphere/library/techarticles/0402\\_du/0402\\_du.html](http://www.ibm.com/developerworks/websphere/library/techarticles/0402_du/0402_du.html).

## Procedure

Run the client once: Run the `MQ_INSTALLATION_PATH \tools\wcf\samples\WAS\default\client\bin\Release\TestClient.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM MQ. The client application starts both of the service methods at the same time, sending two messages to the sample queue.

## Results

The service application gets the messages from the sample queue and provides a response to the `HelloWorld(...)` method call which the client application outputs to the console.

## Problem determination on the WCF custom channel for IBM MQ

You can use IBM MQ trace to collect detailed information about what various parts of the IBM MQ code is doing. When using Windows Communication Foundation (WCF), a separate trace output is generated for the WCF custom channel trace integrated with the Microsoft WCF infrastructure trace.

Fully enabling trace for the WCF custom channel produces two output files:

1. The WCF custom channel trace integrated with the Microsoft WCF infrastructure trace.
2. The WCF custom channel trace integrated with XMS .NET.

By having two trace outputs, problems can be tracked at each interface using the appropriate tools, for example:

- WCF problem determination using suitable Microsoft tooling.
- IBM MQ MQI client issues using the XMS trace format.

To simplify trace enablement, the .NET 3 TraceSource and XMS .NET trace stack are both controlled using a single interface as described in: "WCF trace configuration and trace file names: SOAP/JMS interface" on page 1211.

## WCF custom channel exception hierarchy

The exceptions types thrown by the custom channel are consistent with WCF and are typically a `TimeoutException` or `CommunicationException` (or a subclass of `CommunicationException`). Further details of the error condition, where available, are provided using linked or inner exceptions.

## SOAP/JMS interface

The following exceptions are typical examples, and each layer in the architecture of the channel contributes an additional linked exception, for example `CommunicationsException` has a linked `XMSException`, which has a linked `MQException`:

1. `System.ServiceModel.CommunicationsExceptions`
2. `IBM.XMS.XMSException`
3. `IBM.WMQ.MQException`

Key information is captured and provided in the data collection of the highest `CommunicationException` in the hierarchy. This capture and provision of data prevents the need for the applications to link to each layer in the architecture of the channel in order to interrogate the linked exceptions, and any additional information they might contain. The following key names are defined:

- `IBM.XMS.WCF.ErrorCode`: The error message code of the current custom channel exception.
- `IBM.XMS.ErrorCode`: The error message of the first XMS exception in the stack.
- `IBM.WMQ.ReasonCode`: The underlying IBM MQ reason code.

- `IBM.WMQ.CompletionCode`: The underlying IBM MQ completion code.

## Non-SOAP/Non-JMS interface

The following exceptions are typical examples, and each layer in the architecture of the channel contributes an additional linked exception, for example `CommunicationsException` has a linked `MQException`:

1. `System.ServiceModel.CommunicationsExceptions`
2. `IBM.WMQ.MQException`

Key information is captured and provided in the data collection of the highest `CommunicationException` in the hierarchy. This capture and provision of data prevents the need for the applications to link to each layer in the architecture of the channel in order to interrogate the linked exceptions, and any additional information they might contain. The following key names are defined:

- `IBM.WMQ.WCF.ErrorCode`: The error message code of the current custom channel exception.
- `IBM.WMQ.ReasonCode`: The underlying IBM MQ reason code.
- `IBM.WMQ.CompletionCode`: The underlying IBM MQ completion code.

## WCF trace configuration

There are two options for configuring WCF trace. You can either configure trace programmatically or through an environment variable.

### WCF trace configuration and trace file names: SOAP/JMS interface:

When trace is fully enabled, it produces two output files, one for diagnosing WCF problems, and one detailed file for internal trace diagnostic material. To simplify trace enablement, both the .NET 3 `TraceSource` and XMS .NET trace stacks use a single interface.

Two different trace methods are available for the WCF custom channel. The two trace methods are activated independently or together. Each method produces its own trace file, so when both trace methods have been activated, two trace output files are generated.

To keep the configuration and enablement as simple as possible, the same interface is used to control both trace methods. The `app.config` file must be edited to include the relevant trace configuration as described in the following section. Users can then add their own equivalent sections to combine the output with trace from their own application.

WCF custom channel tracing is not enabled by default. You must first create a trace listener, then set the required trace level for the selected trace source in the `app.config` file.

### Configuring WCF custom channel with WCF infrastructure trace

Add the following section of code to the `<system.diagnostics><sources>` section in the `app.config` file:

```
<source name="IBM.XMS.WCF" switchValue="Verbose,ActivityTracing">
 <listeners>
 <remove name="Default"/>
 <add name="NewListener"/>
 </listeners>
</source>
```

The preceding piece of code makes the channel trace using the .NET 3 `TraceSource`. All invocations of the configuration files associated with the executable files are controlled by this piece of code.

## Configuring WCF custom channel with XMS .NET trace

Configuring the XMS .NET trace requires that you add a section of code to the <system.diagnostics><sources> section in the app.config file. However, the piece of code is added to the extensible <source> element shown in the Configuring WCF custom channel with WCF infrastructure trace section. So although the WCF infrastructure trace code must be present for the XMS .NET trace to work, the WCF infrastructure trace can be disabled if it is not required, as described in the Enabling WCF trace section.

```
<source name="IBM.XMS.WCF" switchValue="Verbose, ActivityTracing"
 xmsTraceSpecification="*=all=enabled" xmsTraceFilePath="path"
 xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced">
 <listeners>
 <remove name="Default"/>
 <add name="NewListener"/>
 </listeners>
</source>
```

## WCF trace configuration variables

Table 177. WCF trace configuration variables

Variable	Description
name	Specify the name as: IBM.XMS.WCF
switchValue	The switchValue controls the trace level. When switchValue is set to Off, the WCF infrastructure TraceSource is not generated. Any other value, such as Verbose, generates TraceSource. For detailed trace level information from Microsoft, consult your WCF documentation or go to the Microsoft WCF Tracing web page: <a href="http://msdn.microsoft.com/en-us/library/ms733025(vs.85).aspx">http://msdn.microsoft.com/en-us/library/ms733025(vs.85).aspx</a>
xmsTraceSpecification= <i>ComponentName</i> = <i>type</i> = <i>state</i>	<p><i>ComponentName</i> is the name of the class that you want to trace. You can use a * wildcard character in this name. For example:</p> <p><code>*=all=enabled</code></p> <p>specifies that you want to trace all classes, and <code>IBM.XMS.impl.*=all=enabled</code></p> <p>specifies that you require API trace only.</p> <p><i>type</i> can be any of the following trace types:</p> <ul style="list-style-type: none"> <li>• all</li> <li>• debug</li> <li>• event</li> <li>• EntryExit</li> </ul> <p><i>state</i> can be either enabled or disabled.</p>
xmsTraceFilePath=" <i>filename</i> "	If you do not specify an xmsTraceFilePath, or if the xmsTraceFilePath is present but contains an empty string, then the trace file is placed in the current directory. To store the trace file in a named directory, specify the directory name in the xmsTraceFilePath, for example: <code>xmsTraceFilePath="c:\somepath"</code>

Table 177. WCF trace configuration variables (continued)

Variable	Description
<code>xmsTraceFileSize="size"</code>	The maximum allowed size of the trace file. When a file reaches this size, it is archived and renamed. The default maximum is 20 KB, which is specified as: <code>xmsTraceFileSize="20000000"</code> .
<code>xmsTraceFileNumber="number"</code>	The number of trace files that are to be retained. The default is 4 (one active file and three archive files). The minimum number allowed is two.
<code>xmsTraceFormat="format"</code>	There are two levels of <code>xmsTraceFormat</code> : basic and advanced. The default trace format is basic if you do not specify an <code>xmsTraceFormat</code> , or if the <code>xmsTraceFormat</code> is present but contains an empty string. Trace files are produced in this format if you specify: <code>xmsTraceFormat="basic"</code>  If you require trace that is compatible with trace analyzer tools, you must specify: <code>traceFormat="advanced"</code>

### Enabling WCF trace

There are four combinations for enabling and disabling the two different trace methods. The four combinations require editing the values of the sections of coded described in the preceding sections.

There is also an environment variable that can be set; for more information see “Enabling WCF trace with the WCF\_TRACE\_ON environment variable” on page 1214.

This table and the values shown are dependent on the pieces of code demonstrated earlier have already been added to the app.config file.

Table 178. WCF trace enablement combinations.

Trace type	Value changed	Example
XMS trace enabled. WCF TraceSource enabled	The switchValue is not set to Off	<pre>&lt;source name="IBM.XMS.WCF" switchValue=" Verbose, ActivityTracing" xmsTraceSpecification="*=all=enabled" xmsTraceFilePath="path" xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced"&gt;   &lt;listeners&gt;     &lt;remove name="Default"/&gt;     &lt;add name="NewListener"/&gt;   &lt;/listeners&gt; &lt;/source&gt;</pre>
XMS trace enabled. WCF TraceSource disabled	The switchValue is set to Off and an xmsTraceSpecification has been given	<pre>&lt;source name="IBM.XMS.WCF" switchValue="Off, ActivityTracing" xmsTraceSpecification="*=all=enabled" xmsTraceFilePath="path" xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced"&gt;   &lt;listeners&gt;     &lt;remove name="Default"/&gt;     &lt;add name="NewListener"/&gt;   &lt;/listeners&gt; &lt;/source&gt;</pre>

Table 178. WCF trace enablement combinations. (continued)

Trace type	Value changed	Example
XMS trace disabled. WCF TraceSource enabled	<p>There are two ways to achieve this result:</p> <ul style="list-style-type: none"> <li>The switchValue variable is not set to Off and an xmsTraceSpecification has not been added</li> <li>The switchValue variable is not set to Off and the xmsTraceSpecification has been set to disabled</li> </ul>	<pre>&lt;source name="IBM.XMS.WCF" switchValue="Verbose, ActivityTracing" xmsTraceSpecification="*=all=disabled" xmsTraceFilePath="path" xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced"&gt; &lt;listeners&gt; &lt;remove name="Default"/&gt; &lt;add name="NewListener"/&gt; &lt;/listeners&gt; &lt;/source&gt;</pre>
XMS trace disabled. WCF TraceSource disabled	<p>There are three ways to achieve this result:</p> <ul style="list-style-type: none"> <li>No &lt;source&gt; element in the app.config file</li> <li>The switchValue variable is set to Off and an xmsTraceSpecification has not been added</li> <li>The switchValue variable is set to Off and the xmsTraceSpecification has been set to disabled</li> </ul>	<pre>&lt;source name="IBM.XMS.WCF" switchValue="Off, ActivityTracing" xmsTraceSpecification="*=all=disabled" xmsTraceFilePath="path" xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced"&gt; &lt;listeners&gt; &lt;remove name="Default"/&gt; &lt;add name="NewListener"/&gt; &lt;/listeners&gt; &lt;/source&gt;</pre>

### Enabling WCF trace with the WCF\_TRACE\_ON environment variable

As well as the preceding methods described to enable WCF trace, the XMS .NET trace can also be enabled using the WCF\_TRACE\_ON environment variable.

Setting the WCF\_TRACE\_ON environment variable to any non-null value is the equivalent of setting the xmstraceSpecification to \*=all=enabled, for example: " set WCF\_TRACE\_ON=true "

However, if the xmstraceSpecification is explicitly set in the app.config file, then the WCF\_TRACE\_ON environment variable is overridden.

### WCF trace output files and file names

XMS trace files are traditionally named using the base name and process ID format of: xms\_trace\_pid.log, where pid is the process ID.

As XMS trace files can still be produced in parallel with WCF custom channel trace files, the WCF custom channel trace integrated with XMS .NET trace output files have the following format to avoid confusion: wcf xms\_trace\_pid.log, where pid is the process ID.

The trace output file is created in the current working directory by default, but this destination can be redefined if necessary.

## WCF trace configuration: Non\_SOAP/Non-JMS interface:

For the Non\_SOAP/Non-JMS interface, you can either configure trace through an environment variable or programmatically.

There are two ways of enabling trace for the Non-SOAP/Non-JMS interface:

- By setting WMQ\_TRACE\_ON as the environment variable.
- By adding the following section of code to the <system.diagnostics><sources> section in the app.config file

```
<source name="IBM.WMQ.WCF" switchValue="Verbose, ActivityTracing"
 xmsTraceSpecification="*=all=enabled"
 xmsTraceFileSize="2000000" xmsTraceFileNumber="4"
 xmsTraceFormat="advanced">
</source>
```

## WCF XMS First Failure Support Technology ( FFST )

You can collect detailed information about what various parts of the IBM MQ code is doing by using IBM MQ trace. XMS FFST has its own configuration and output files for the WCF custom channel.

XMS FFST trace files are traditionally named using the base name and process ID format of: *xmsffdc pid\_date.txt*, where *pid* is the process ID and *date* is the time and date.

As XMS FFST trace files can still be produced in parallel with WCF custom channel XMS FFST files, the WCF custom channel XMS FFST output files have the following format to avoid confusion: *wcf ffdc pid\_date.txt*, where *pid* is the process ID and *date* is the time and date.

This trace output file is created in the current working directory by default, but this destination can be redefined if necessary.

The WCF custom channel with XMS .NET trace header is similar to the following example:

```
***** Start Display XMS WCF Environment *****
Product Name :- value
WCF Version :- value
Level :- value
***** End Display XMS WCF Environment *****
```

The FFST trace files are formatted in the standard way, without any formatting that is specific to the custom channel.

## WCF version information

WCF version information aids with problem determination and is included in the assembly metadata of the custom channel.

The IBM MQ custom channel for WCF version metadata can be retrieved in one of three ways:

- Using the IBM MQ utility dspmqver. For information about how to use dspmqver, see: dspmqver
- Using the Windows Explorer properties dialog: In the Windows Explorer, right-click **IBM.XMS.WCF.dll** > **Properties** > **Version**.
- From the header information of any of the channels FFST or trace files. For more information about the FFST header information, see: "WCF XMS First Failure Support Technology ( FFST )"

## WCF hints and tips

The following hints and tips are in no significant order, and might be added to when new versions of the documentation are released. They are subjects that might save you time if they are relevant to the work that you are doing.

### Externalizing exceptions from the WCF service host:

For services hosted using the WCF service host; any unhandled exceptions thrown by the service, WCF internals, or channel stack are not externalized by default. To be informed of these exceptions, an error handler must be registered.

The following code provides an example of defining the error handler service behavior which can be applied as an attribute of a service:

```
using System.ServiceModel.Dispatcher;
using System.Collections.ObjectModel;
....
public class ErrorHandlerBehaviorAttribute : Attribute, IServiceBehavior, IErrorHandler
{
 //
 // IServiceBehavior Interface
 //
 public void AddBindingParameters(ServiceDescription serviceDescription,
 ServiceHostBase serviceHostBase, CollectionServiceEndpoint endpoints,
 BindingParameterCollection bindingParameters)
 {
 }
 public void ApplyDispatchBehavior(ServiceDescription serviceDescription,
 ServiceHostBase serviceHostBase)
 {
 foreach (ChannelDispatcher channelDispatcher in serviceHostBase.ChannelDispatchers) {
 channelDispatcher.ErrorHandlers.Add(this);
 }
 }
 public void Validate(ServiceDescription serviceDescription, ServiceHostBase serviceHostBase)
 {
 }

 //
 // IErrorHandler Interface
 //
 public bool HandleError(Exception e)
 {
 // Process the exception in the required way, in this case just outputting to the console
 Console.Out.WriteLine(e);

 // Always return false to allow any other error handlers to run
 return false;
 }
 public void ProvideFault(Exception error, MessageVersion version, ref Message fault)
 {
 }
}
```



## Handling different SOAP response element names:

WCF expects the name of a returned value to be in a specific format by default, but a service might not return an element with its name in the expected format.

WCF has the convention of expecting the returned value to be named in the following format: *methodName* Result where *methodName* is the name of the service operation. For example, for a service called `getQuote`, WCF expects the response to be called: `getQuoteResult`.

However, the service can return an element with a name that does not conform to this format.

When running the `scvutil` tool to generate a proxy client, if the WSDL specifies a different name, then the proxy interface adds parameters to instruct WCF with the name to look for. For example:

```
[System.ServiceModel.OperationContractAttribute(Action = "", ReplyAction = "*")]
[System.ServiceModel.XmlSerializerFormatAttribute(Style = System.ServiceModel.OperationFormatStyle.Rpc,
 Use = System.ServiceModel.OperationFormatUse.Encoded)]
[return: System.ServiceModel.MessageParameterAttribute(Name = "getQuoteReturn")]
float getQuote(string in0);
```

If you create your own interface (for example, by adding a request-reply method to an existing proxy interface), then you must ensure that you add the same parameters to the interface if the service returns a different name. If you do not do so, then the most common issue is that a call to the service method always returns a null value; if an object is returned, then the method returns null, but if a numeric value such as an integer is returned, then the method returns 0.

---

## Developing web services with IBM MQ

You can develop IBM MQ applications for web services using the IBM MQ transport for SOAP.

### About this task

**Note:** From Version 9.0, the IBM MQ transport for SOAP is deprecated. This includes the following product features:

- IBM MQ Java Listener
- IBM MQ .NET 1 and 2 Listener
- IBM MQ Java Axis2 Client
- IBM MQ Java client (deprecation announced in Version 8.0)
- IBM MQ .NET 1 and 2 clients (deprecation announced in Version 8.0)
- IBM MQ bridge for HTTP (deprecation announced in Version 8.0)

The IBM MQ transport for SOAP provides a JMS transport for SOAP. The IBM MQ transport for SOAP is also integrated into other environments such as Microsoft Windows Communication Foundation, WebSphere Application Server, and CICS Transaction Server.

For more information about the IBM MQ transport for SOAP, see “Developing web services with IBM MQ transport for SOAP” on page 1218.

**Related concepts:**

“Application development concepts” on page 2

You can use a choice of procedural or object-oriented languages to write IBM MQ applications. Use the links in this topic for information about IBM MQ concepts that are useful for application developers.

“Developing applications for IBM MQ” on page 1

You can develop applications to send and receive messages, and to manage your queue managers and related resources. IBM MQ support applications written in procedural languages, and object oriented languages and frameworks.

“Design considerations for IBM MQ applications” on page 43

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by IBM MQ.

“Writing a procedural application for queuing” on page 74

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

“Writing client procedural applications” on page 285

What you need to know to write client applications on IBM MQ using a procedural language.

“Writing publish/subscribe applications” on page 172

Start writing publish/subscribe IBM MQ applications.

“Building a procedural application” on page 382

You can write an IBM MQ application in one of several procedural languages, and run the application on several different platforms.

“Handling procedural program errors” on page 436

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

**Related tasks:**

“Using the IBM MQ sample procedural programs” on page 457

These sample programs are written in procedural languages, and demonstrate typical uses of the Message Queue Interface (MQI). IBM MQ programs on different platforms.

## Developing web services with IBM MQ transport for SOAP

IBM MQ transport for SOAP provides a JMS transport for SOAP. The IBM MQ transport for SOAP is also integrated into other environments such as Microsoft Windows Communication Foundation, WebSphere Application Server, and CICS Transaction Server.

### About this task

**Note:** From Version 9.0, the IBM MQ transport for SOAP is deprecated. This includes the following product features:

- IBM MQ Java Listener
- IBM MQ .NET 1 and 2 Listener
- IBM MQ Java Axis2 Client
- IBM MQ Java client (deprecation announced in Version 8.0)
- IBM MQ .NET 1 and 2 clients (deprecation announced in Version 8.0)
- IBM MQ bridge for HTTP (deprecation announced in Version 8.0)

## Introduction to IBM MQ transport for SOAP

The IBM MQ transport for SOAP provides a JMS transport for SOAP. The IBM MQ SOAP sender and listener provide a means to call web services.

The IBM MQ SOAP listener supports services hosted by .NET Framework 1, .NET Framework 2, and Axis 1.4. The IBM MQ SOAP sender supports web services clients running on .NET Framework 1, .NET Framework 2, Axis 1.4 and Axis2. Clients can be either an IBM MQ server or client application. The IBM MQ transport for SOAP is also integrated into other environments such as Microsoft Windows Communication Foundation, WebSphere Application Server, and CICS Transaction Server.

The integration into Microsoft Windows Communication Foundation is part of the IBM MQ support for .NET Framework 3.

The IBM MQ transport for SOAP is a set of protocols and tools to transport SOAP messages by using JMS over IBM MQ. It is packaged in different ways for different application environments as shown in Table 179.

Table 179. IBM MQ transport for SOAP application environments

	Integrated with additional IBM MQ components	Integrated into a framework
Provided as part of IBM MQ installation	.NET Framework 1 .NET Framework 2 Axis 1.4	Windows Communication Foundation (.NET Framework 3) Axis2 (Client only)
Provided in another software package		WebSphere Application Server CICS Transaction Server 4.1 WebSphere ESB WebSphere Process Server for Multiplatforms

Integration of the IBM MQ transport for SOAP into an application framework simplifies the development and deployment of web services to IBM MQ.

With additional IBM MQ SOAP components, you can interact directly with the IBM MQ SOAP components to develop and deploy services. Use the IBM MQ SOAP tools to configure and deploy the web services and the web service clients to IBM MQ.

In the integrated environments, development and deployment is simpler. You use the same tools for development and deployment as you would to develop and deploy a SOAP HTTP web service. You must still configure the IBM MQ queues, channels, and queue managers that you require using IBM MQ tools.

You can mix and match IBM MQ SOAP clients and servers from any of these environments.

### Benefits

The IBM MQ transport for SOAP offers existing IBM MQ users the following principal benefits:

#### Using your IBM MQ network to connect existing web services.

The services might be ones you have written, or services that are provided as interfaces to other packaged software applications you have deployed.

The benefit comes from using your existing IBM MQ network to connect web services. The IBM MQ transport has the advantage of being a managed and reliable queued messaging service.

#### Writing new applications, or converting existing applications, to use SOAP rather than IBM MQ interfaces.

Typically, applications require a specific IBM MQ adapter to be developed to integrate with another application. Adapters have two parts: the connector piece, that puts and gets messages to and from the transport, and the adapter piece that converts data to and from application-specific formats. Integrating each pair of applications is a new challenge.

The benefit of SOAP comes from standardizing on SOAP for defining application interfaces, and then having a choice of transports. You do not need to write application-specific adapters, and you can choose whether to use IBM MQ or HTTP as the connector. Which transport you choose depends on what qualities of service and connectivity you require.

For existing SOAP over HTTP users, the benefit of IBM MQ transport for SOAP comes from using a managed and reliable asynchronous transport. The benefits are twofold:

**A truly asynchronous programming model for availability and performance.**

By using an asynchronous client interface, the client and service applications do not have to be available at the same time. Requests sent by the client will be stored until the service is available to process them.

**A ready built managed network that is designed to be reliable and available.**

By choosing IBM MQ as a transport, you are getting the advantage of using a managed network that provides reliable messaging.

In contrast, transports such as HTTP and FTP over TCP/IP are unmanaged. An unmanaged network is ideal for unpredictable connections: there are fewer management tasks.

## Summary

IBM MQ transport for SOAP provides the following components:

- The SOAP/JMS transport binding is used in WSDL documents to bind a SOAP service to a JMS transport. The IBM MQ implementation of the SOAP/JMS binding uses a URI that takes either of two forms:

**IBM MQ transport for SOAP**

`jms:/queue? &Name=Value&Name=Value...`

**IBM MQ wire format for W3C candidate recommendation**

`jms:queue: qName ?connectionFactory=connectQueueManager (qMgrName)&Name=Value&Name=Value...`

- The mapping of a SOAP message onto an IBM MQ message.
- Two IBM MQ SOAP listeners to receive SOAP requests, one for Java and one for .NET Framework 1 or .NET Framework 2. The listeners use .NET or Axis 1.4 to process the SOAP request.
- Two IBM MQ SOAP senders to create IBM MQ SOAP requests. web services clients register with a sender to process jms: SOAP requests.
- Integration with Windows Communication Foundation (WCF), sometimes known as .NET 3, to send and receive IBM MQ Transport for SOAP messages.
- Integration of the client with Axis2, sometimes known as JAX-WS, to send IBM MQ Transport for SOAP or W3C SOAP JMS messages.
- The command **amqwdployMQService**, which creates development and run time components and scripts to deploy a web service by using the IBM MQ transport for SOAP.
- Sample Java and .NET client and service code.
- A script to set the class path, and other utility scripts.

In the integrated environments the sender and listener are integrated into each environment, as are extensions to the development and deployment tools.

## Integration of SOAP and IBM MQ:

The IBM MQ transport for SOAP extends SOAP, and web services tools and run time, with IBM MQ as an alternative transport to HTTP for SOAP. You do not need to modify existing web services to use IBM MQ transport for SOAP as a transport. The transport uses a custom URI format for SOAP/JMS. The W3C URI format for SOAP/JMS is supported in a limited way by Axis2 clients.

An additional line of code has to be added to clients in the .NET Framework 1, .NET Framework 2, and Axis 1.4 environments. No additional code is required in Axis 2 and Windows Communication Foundation (WCF) clients. The IBM MQ SOAP listener runs services in the .NET Framework 1, .NET Framework 2, and Axis 1.4 environments. The IBM MQ transport for SOAP is integrated into some other application server environments, including WCF, CICS, and WebSphere Application Server.

### What is SOAP?

SOAP<sup>9</sup> describes the standardized format of the messages and interaction protocols that applications use to exchange requests, replies, and datagrams. SOAP is independent of the transport used to transfer the messages, and of the application environment that sends and receives the messages. The W3C defines SOAP Version 1.2 succinctly:

*SOAP Version 1.2 provides the definition of the XML-based information which can be used for exchanging structured and typed information between peers in a decentralized, distributed environment.*<sup>10</sup>

To use SOAP it must be bound to a transport, such as HTTP, e-mail, or IBM MQ.

A SOAP protocol binding framework is the set of rules for carrying a SOAP message on top of another protocol, such as HTTP. SOAP Version 1.2 Part 2: Adjuncts (Second Edition) describes the SOAP HTTP binding.

The W3C candidate recommendation, 4 June 2009, SOAP over Java Message Service 1.0, describes the recommendation for the SOAP JMS binding. As JMS is an API specification, and not a transport protocol, the JMS SOAP recommendation does not describe the wire-format of SOAP JMS messages. It describes the SOAP interaction protocols and the JMS API binding. Consequently, when using the JMS SOAP recommendation you must still use the same JMS implementation for the SOAP client and the SOAP server. It does enable a SOAP JMS application to be run on any implementation of JMS. A JMS implementation can be plugged into a J2EE application server, if both the server and the JMS implementation comply with the JCA specification. IBM MQ JMS complies with the JCA specification and can be plugged into a compliant application server.

IBM MQ transport for SOAP binding is like the proposed W3C standard, but it is not the same. Its usage is described in the topic MQRFH2 SOAP settings. Unlike the W3C candidate recommendation, the SOAP binding is not formally specified. Effectively, it is the HTTP binding, and the service address takes the form, `jms:/queue?name=value&name=value...`, rather than `http://authority/path?query#fragment`. `jms:` is not an officially registered IANA URI scheme.

### What is a web service?

SOAP enables programs written in different languages, running on different platforms, to communicate using various transport protocols. SOAP is the protocol specification. A web service is an application that provides a service through a SOAP interface that can be accessed using internet protocols.

---

9. Historically, the acronym stood for Simple Object Access Protocol.

10. W3C: SOAP Version 1.2 Part 0

An important goal of SOAP is to provide services that clients can use easily. Once you have designed a client to use a service, you can program the call to invoke the service without reference to external documentation. Service interfaces are described in XML, in a WSDL document. The query, `http://authority/path?wsdl`, returns the WSDL description of a SOAP service.

**Tip:** When you deploy a web service to use IBM MQ, also deploy the service to HTTP so that the standard WSDL query works.

## Developing web services

Web services have a client and a service part. The service is written first, either starting from the interface description in WSDL, or by following the rules for writing the service class. Web service toolkits have utilities to generate WSDL from the interface definition of a class; for example `java2wsdl` or `disco`. They also have tools to generate or class skeletons from WSDL interface descriptions; for example `wsdl2java`, `wsimport`, or `wsdl`. The former is known as bottom up development, and the latter top down.

The `amqwdeployWMQService` command in IBM MQ transport for SOAP uses these tools to generate WSDL, client stubs, and client proxies.

Web services are typically written using an integrated development environment targeted at a particular application server environment:

### Eclipse IDE for Java EE Developers

Creates web services for Axis 2. Supports JAX-RPC and JAX-WS

### Rational® Application Developer V7.5

Creates web services for WebSphere Application Server V7 and previous versions, and also for Axis. Supports JAX-RPC and JAX-WS.

### WebSphere Integration Developer V6.2

Creates web services for WebSphere Process Server and WebSphere ESB. Supports JAX-RPC and JAX-WS.

### Visual Studio 2012

Creates web services for .NET Framework 3.5 and earlier ( Windows Communication Foundation)

You can use any of these tools in combination with IBM MQ transport for SOAP. Once you have developed a service to use with HTTP, use the `amqwdeployWMQService` tool to deploy the services to use IBM MQ as a transport. You can write a new client using the output from the tool, or modify your existing clients to use the IBM MQ transport for SOAP.

If IBM MQ transport for SOAP is integrated into the application environment, then you do not need use the `amqwdeployWMQService` tool or modify the client code. The client SOAP layer directs client requests that have a URI with the prefix `jms:` to the IBM MQ transport for SOAP. The server SOAP layer calls IBM MQ transport for SOAP to wait for `jms:` SOAP requests, and returns responses to IBM MQ transport for SOAP.

Typically, .NET services have been developed bottom up using web service annotations in code, and Java services top down, using WSDL interface definitions. The difference in approaches is narrowing, as Java Standard Edition Version 6 supports JAX-WS 2.0, and uses annotations to qualify the definition of service interfaces. It is now as easy to develop Java services bottom up as top down. Which approach you choose is a matter of development method.

The web services client is written after the service, using the WSDL service definition and generated client stubs and proxies. In some applications, the service definition is not known when the client is written. The client retrieves the service WSDL and creates service requests dynamically. More commonly the service definition is known, but the address to which the service is deployed, is not. The web service toolkit generates interfaces for the client to use to make service requests. The client provides the service

address when it is required. In the third case, the WSDL contains all the information a client needs. The WSDL contains both the interface and address of the service. The code generated by the web service toolkit has all the information needed by the client to make requests of a service.

You can use any of these three styles with IBM MQ transport for SOAP.

### Web service application environments

Web service toolkits require a mapping from the WSDL definition of a service to the streams of bytes that are transferred in SOAP requests and responses. The byte stream is defined by the SOAP specification, and is contained in the SOAP envelope. The SOAP envelope is shown in Figure 152.

```
<?xml version='1.0'?>
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
<soap:Header> <!-- optional -->

<!-- headers... -->

</soap:Header>
<soap:Body>
<!-- payload or fault message -->

</soap:Body>
</soap:Envelope>
```

Figure 152. SOAP envelope

Mapping from the SOAP envelope to the language binding and back is part standardized, and part proprietary. The mapping is fundamental to the .NET architecture, and is provided as part of the Common Language Runtime (CLR). The mapping is standardized in Java by JAX specifications. Because the Java mappings are standardized, Java web service clients and services are portable between different Java-based application environments. JAX-RPC (sometimes called JAX-WS 1.0) is the mapping most in use today. It is supported by Axis 1.4. JAX-WS (sometimes called JAX-WS 2.0) is a greatly improved standard and is likely to replace JAX-RPC rapidly. JAX-WS is supported by Axis 2.0. IBM MQ 7.0.1 does not support JAX-WS and Axis 2.

IBM MQ transport for SOAP does not alter the contents of the SOAP envelope, and the contents do not affect the transport. The language bindings do affect the IBM MQ transport for SOAP. IBM MQ 7.0.1 supports .NET Framework 1, .NET Framework 2, and Axis 1.4 using the code and utilities shipped with IBM MQ transport for SOAP. Support for the WebSphere transport for SOAP in .NET Framework 3 and 3.5 is implemented using the IBM MQ custom channel for Windows Communication Foundation.

Other SOAP development and run time environments might ship support for IBM MQ transport for SOAP, and support different languages. For example, web services running on CICS supports languages such as COBOL and PL/1.

**Note:** The mapping used makes no difference to the interoperability of web services. You can mix and match clients and services written using .NET, JAX-RPC, and JAX-WS mappings.

### What is IBM MQ transport for SOAP?

IBM MQ transport for SOAP is a SOAP binding and a web services toolkit. Together, they enable applications to exchange SOAP messages using IBM MQ rather than HTTP. Figure 153 on page 1224 shows IBM MQ as an alternative to HTTP as a SOAP transport.




Figure 153. Overview of IBM MQ transport for SOAP

SOAP over HTTP is shown as (1) in the diagram. The client SOAP layer converts a request into a SOAP message and the HTTP component sends over TCP/IP. The HTTP server component listens for HTTP requests, typically on the TCP/IP port 80. If the request is for a SOAP service, the HTTP server component calls the SOAP layer to convert the SOAP request into method call. It then returns the response.

SOAP over IBM MQ is shown as (2). The client application registers the IBM MQ SOAP sender component as a handler for the `jms:` protocol with the SOAP layer. The SOAP layer passes SOAP messages addressed to `jms:` to the IBM MQ SOAP sender. The sender uses the URI in the message to place the message on the request queue with the required qualities of service. The corresponding IBM MQ SOAP listener waits for messages on its request queue and calls the SOAP layer to process requests and return responses.

The SOAP sender and listener are normal IBM MQ programs. They can be connected to the same queue manager, as in Figure 154 on page 1225, or connected to different queue managers; see Figure 155 on page 1226. The client can be connected by a client connection.



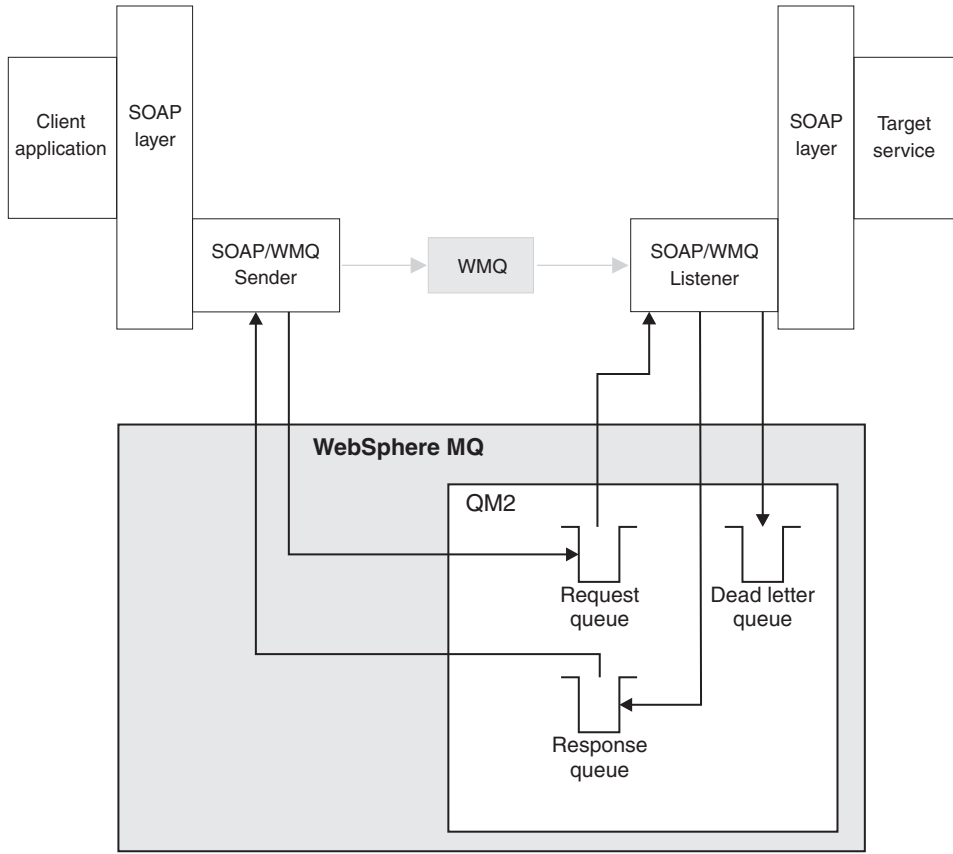


Figure 154. Queues used by SOAP/IBM MQ (single queue manager)

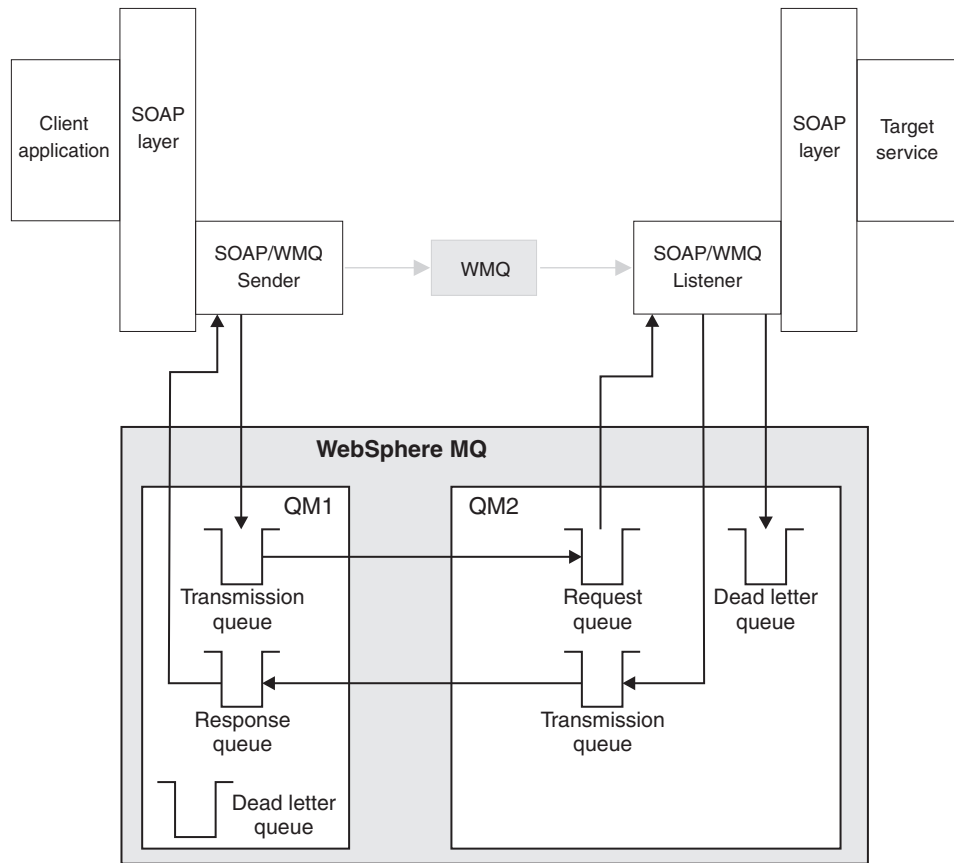


Figure 155. Queues used by SOAP/IBM MQ (separate queue managers)

### W3C candidate recommendation for binding SOAP to JMS.

The W3C candidate recommendation defines the SOAP over JMS binding; SOAP over Java Message Service 1.0. Also useful for its examples is URI Scheme for Java(tm) Message Service 1.0<sup>11</sup>.

Some application frameworks, such as WebSphere Application Server v7, have support for the W3C candidate recommendation. Send SOAP requests formatted with a URI compatible with the W3C candidate recommendation using the Axis2 client; see W3C SOAP over JMS URI for the IBM MQ Axis 2 client. The Axis2 client sends a SOAP request formatted with either a W3C or an IBM MQ transport for SOAP based on URI in the SOAP request.

Axis2 client support for the W3C recommendation is introduced in the 7.0.1.3 fix pack. Support for other clients, and for the SOAP listeners provided by IBM MQ is not provided.

11. Look for *URI Scheme for JMS*, in the W3C specification references, for the latest draft.

## Implementation of WebSphere transport for SOAP on .NET Framework 1, .NET 2 and Axis 1.4:

You might want to write your own IBM MQ SOAP sender and listener. Use the implementation of IBM MQ transport for SOAP on .NET Framework 1, .NET Framework 2, and Axis 1.4 as a guide.

1. A client program uses the appropriate web services framework in the same way as it would for the HTTP transport. It must also register the `jms:` prefix. The prefix is registered using either the `com.ibm.mq.soap.Register.extension()` Java method or the `IBM.WMQSOAP.Register.Extension()` CLR method.
2. The Axis 1.4 or .NET Framework 1 or 2 framework marshals the call into a SOAP request message exactly as for SOAP/HTTP.
3. An IBM MQ service is identified by a URI prefixed with `jms:`. When the framework identifies the `jms:` URI, it calls the IBM MQ transport sender code; `com.ibm.mq.soap.transport.jms.WMQSender` (for Axis 1.4) or `IBM.WMQSOAP.MQWebRequest` (for .NET1 and 2). If the framework encounters a URI with an `http:` prefix, it calls the standard SOAP over HTTP sender.
4. The SOAP message is transported by the IBM MQ SOAP sender using the request queue. The **SimpleJavaListener** (for Java ) or **amqwSOAPNETListener** (for .NET) receives the request message. The IBM MQ SOAP listeners are stand-alone processes and are multithreaded with a customizable number of threads.
5. The IBM MQ SOAP listener reads the incoming SOAP request, and passes it to the appropriate web service infrastructure.
6. The web service infrastructure parses the SOAP request message and invokes the service. The procedure is the same as for a message that arrived on an HTTP transport.
7. The infrastructure formats the response into a SOAP response message and returns it to the IBM MQ SOAP listener.
8. The listener places the message on the response queue, and the message is transferred to the IBM MQ SOAP sender. The sender passes it to the client web service infrastructure.
9. The client infrastructure parses the response SOAP message and hands the result back to the client application.

Each application context is served by a separate IBM MQ request queue.

The application context is controlled in Axis 1.4 by ensuring that the IBM MQ SOAP listener and service execute in the appropriate directory. Axis 1.4 sets the correct CLASSPATH for the directory.

Application context is controlled in .NET by the IBM MQ SOAP listener executing the service in a context created by a call to `ApplicationHost.CreateApplicationHost`. The call specifies the target execution directory. Each service then operates in the directory in which it was deployed.

**amqwdeployWmqService** generates the request and response queues. It also generates the infrastructure necessary for handling the queues and deploying services to Axis 1.4.

## IBM MQ transport for SOAP and web services reliable messaging:

Web services reliable messaging is a protocol for reliably exchanging web service requests and responses over an unreliable connection. It is best-suited to solve problems of short-lived connection disruption.

IBM MQ for SOAP takes advantage of using an IBM MQ managed and reliable network for passing SOAP messages. Transports such as HTTP and FTP are unmanaged. Unmanaged networks are ideal for unpredictable connections, where the difficulties and costs of managing connections outweigh the benefits of not losing requests and responses.

To overcome the problem of losing files when connections break in unmanaged networks, services like managed FTP build a management layer on top of FTP. The management layer takes over the burden of checking that files have transferred successfully from users, and retransmits missing files if necessary. To use managed FTP, you must have the management software installed at both ends of the connection.

Web services reliable messaging (WSRM) takes a different approach to solving the problem of unreliable connections. Its goal is to transfer web service requests and responses reliably, without both ends of the connection having to use the same software. Any software, by implementing the web services reliable messaging protocol, can exchange messages reliably with other software.

When a connection fails, a sender and receiver must preserve the context of the WSRM message transfer, using a generated URI as a key. The sender and receiver keep attempting to establish a new connection. If a new connection is successfully established, the transfer completes. The WSRM specification does not specify how context is preserved, or when a new connection is attempted.

You might decide that only short-lived outages are of interest. For longer outages, you might be prepared to discard transfers that could not be restarted after a time. Similarly, you might be prepared to discard transfers if either the client or service fails. Leaving the user responsible for assuring transfers, places fewer demands on managing the coordination of client and service.

If network outages are long-lived, over 30 minutes or so, or if the client or server fails, there is an increased likelihood that some connections are never re-established. You can no longer rely on WSRM restoring message transfer automatically in an unmanaged way. You have to consider managing the failed WSRM connections, which implies developing software to manage the network of clients and services.

Using WSRM to overcome short outages might significantly reduce dealing with lost messages on a mobile network. If you do not have to assure message delivery, the benefits of reducing message loss might justify the additional cost of developing a WSRM implementation.

SOAP over JMS provides assured message delivery and deals with longer duration outages of the client, the server, and the network. If you are seeking a more reliable quality of service for SOAP than HTTP, which solution do you choose: IBM MQ transport for SOAP or WSRM? The answer depends on many factors. Some of the factors to consider are listed:

1. Whether the unreliability is due to connection failure.
2. How long-lived connection failures are.
3. If you can manage both the client and server side of the connection.
4. If the user or an administrator is ultimately responsible for the delivery of messages.

## Installing and verifying IBM MQ web services

Use the instructions in these topics to install and verify the IBM MQ transport for SOAP.

### Installing IBM MQ Web transport for SOAP:

Use these instructions to install the IBM MQ Web transport for SOAP. The installation creates tools to run web service clients or services using IBM MQ as the SOAP transport. The tools are used in the .NET Framework 1, .NET 2, Axis 1.4, or Axis2 SOAP environments.

### Before you begin

Check the prerequisite products at System Requirements for IBM MQ. The installation process does not check for the presence and availability of the prerequisite software. You must verify that the prerequisites are installed.

IBM MQ provides a copy of the Axis 1.4 run time. Use this version with IBM MQ rather than any other you might have installed. IBM does not provide technical support for Apache Axis. Contact the Apache Software Foundation if you have technical problems with it.

To run web services in the .NET Framework 3 SOAP environment, IBM MQ uses Windows Communication Foundation. The IBM MQ custom channel for Windows Communication Foundation runs web service clients and services using IBM MQ as a transport for SOAP messages.

### About this task

You can install the IBM MQ Web transport for SOAP as either an IBM MQ MQI client or Server application. If you have already installed IBM MQ as a client or a server on your computer, check that you have installed the components listed.

*MQ\_INSTALLATION\_PATH* represents the high-level directory in which IBM MQ is installed.

Carry out the following installation steps.

### Procedure

1. Select the Java and .Net Messaging and web services component for installation.
2. On Solaris and HP-UX, select the Java Runtime environment component for installation.
3. Select the Development toolkit for installation.
4. Install and verify IBM MQ as described in the Quick Beginning for your platform.
5. Copy the Apache Axis 1.4 run time, *axis.jar* from the *prereqs/axis* directory on the IBM MQ installation media. Copy it to the installation directory described in Table 180 on page 1230, Table 181 on page 1230, or Table 182 on page 1230.

#### Windows

```
Copy D:\PreReqs\axis\axis.jar MQ_INSTALLATION_PATH\java\lib\soap
```

#### AIX

```
cp -f PreReqs/axis/axis.jar MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
chown mqm:mqm MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
chmod 444 MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
```

#### HP-UX, Solaris, and Linux (all platforms) installation directories

```
cp -f PreReqs/axis/axis.jar MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
chown mqm:mqm MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
chmod 444 MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
```

6. On Windows 2003, run **Aspnet\_regiis.exe** to update the script maps to point to the version of the Common Language Run time you are using. Look for the **Aspnet\_regiis.exe** utility in `%SystemRoot%\Microsoft.NET\Framework\version-number`.

- Set the environment variable, `WMQSOAP_HOME`, to point to the IBM MQ installation directory.

## Results

Table 180. Windows installation directories

Location	Contents
<code>MQ_INSTALLATION_PATH \programs\bin</code>	Binary, command, DLL, and executable files
<code>MQ_INSTALLATION_PATH \programs\java\lib</code>	.jar files
<code>MQ_INSTALLATION_PATH \programs\java\lib\soap</code>	SOAP .jar files
<code>MQ_INSTALLATION_PATH \programs\soap\samples</code>	Samples and IVT

Table 181. AIX installation directories

Location	Contents
<code>MQ_INSTALLATION_PATH/bin</code>	Shell scripts
<code>MQ_INSTALLATION_PATH/java/lib</code>	.jar files
<code>MQ_INSTALLATION_PATH/java/lib/soap</code>	axis.jar and other JAX-RPC .jar files
<code>MQ_INSTALLATION_PATH/samp/soap</code>	Samples and IVT

Table 182. HP-UX, Solaris, and Linux (all platforms) installation directories

Location	Contents
<code>MQ_INSTALLATION_PATH/bin</code>	Shell scripts
<code>MQ_INSTALLATION_PATH/java/lib</code>	.jar files
<code>MQ_INSTALLATION_PATH/java/lib/soap</code>	axis.jar and other JAX-RPC .jar files
<code>MQ_INSTALLATION_PATH/samp/soap</code>	Samples and IVT

## What to do next

- For .NET only, you must register the IBM MQ transport for SOAP files with the Global Assembly Cache. If .NET is already installed when you install IBM MQ, registration is performed automatically at installation. If you install .NET after IBM MQ, the registration is performed automatically when the IVT is first run.  
You can run `amqregisterdotnet.cmd` to perform registration of the .NET assemblies. You can also run `amqregisterdotnet.cmd` to force reregistration at any stage. Once made, this registration survives system restarts and subsequent reregistration is not normally necessary.
- Run the Installation Verification Test, as described in “Verifying the IBM MQ transport for SOAP” on page 1231.
- If you intend to develop Axis2 client, you must download Axis2 1.4.1 from Apache; see “Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 1249.

## Verifying the IBM MQ transport for SOAP:

Verify the IBM MQ transport for SOAP using the **runivt** command. The command runs a number of demonstration applications and ensures that the environment is correctly set up after installation. Note that the web services part of the transport for SOAP is deprecated, and, if you are a new user, you should not use this.

### Before you begin

Before you run the **runivt** command, ensure that you have the following runtime environments:

- To run on Axis only: you must have a Java SDK (within SOE) available on your system. You must also include the location of the `java.exe` and `javac.exe` commands in the systems **PATH** environment variable.
- To run a test on .NET only (supported only on Windows ): you must have both a Java SDK and the .NET compilers and tools on your system. To do so, access either a Visual Studio command prompt or the Microsoft Windows SDK command prompt, then add the location of the `java.exe` and `javac.exe` files to the **PATH** environment variable.
- To run all available the tests: for Windows platforms, the environment must be configured as described in the .NET test run. On UNIX and Linux platforms, the environment must be configured as described in the Axis only test run.

### About this task

Instead of running the verification test on both .NET and Axis, you might want to run the test only on Axis, or only on .NET.

If you experience problems with the tests and want to start again:

1. Stop the queue manager `WMQSOAP.DEMO.QM` using the `immediate` option.
2. Stop the listener that has been started in a different window.
3. Delete the queue manager.
4. Delete the temporary samples directory you created and start again.

On UNIX and Linux platforms, you must run the command using an X Windows system session.

The **runivt** command changes the contents of the `soap/samples` directory. To keep the installation image unchanged, copy the samples directory to a temporary location, and run the verification test from the temporary location.

You can run the installation verification as many times as you like.

Carry out the following steps to verify the installation of IBM MQ transport for SOAP on .NET Framework 1, .NET Framework 2, and Axis 1.4:

### Procedure

1. Copy the `./tools/soap/samples` directory tree to a temporary location.
2. Start a command window with the temporary directory as the current directory.
3. Use the **runivt** command to start the installation test. The `runivt` script compiles a test class, the sample client, and services before deploying and running them. For the test class, the sample client, and services to run, complete the installation steps outlined in *Installing WebSphere(r) MQ Web transport for SOAP* and ensure that the command prompt used to run the `runivt` command has the required runtime environment set. Use any of the following methods to run the **runivt** command:
  - Run a test on Axis only: `runivt Axis`.
  - Run a test on .NET only (supported only on Windows ): `runivt DotNet`.

- Run all available the tests: runivt.

For more information about the runivt command syntax and parameters, see **runivt: IBM MQ transport for SOAP installation verification test**. The tests that you can run are listed in the file `ivttests.txt` on Windows and `ivttests_unix.txt` on UNIX and Linux platforms.

#### Related information:

runivt: IBM MQ transport for SOAP installation verification test

## Developing web services for IBM MQ transport for SOAP

Use your normal web service development environment to develop services for use with the IBM MQ transport for SOAP.

### Before you begin

1. If you are planning to use the command-line tools supplied with IBM MQ transport for SOAP:
  - a. Create a deployment directory for the service.
  - b. Start a command window in the directory.
  - c. For .NET, `csc.exe` and `wSDL.exe` must be in the path, and be from the same version of .NET Framework.
  - d. For Java,
    - 1) Run the `amqwsetcp` command to set up the classpath.
    - 2) An IBM JRE and a JDK at the same version level must be in the current path. The version level must be at least 5.0.
    - 3) Customize the classpath to include the locations of any additional `.jar` libraries, and directories containing `.java` packages, including for the service you are developing. Put the current directory `“.”` in the classpath.
    - 4) Create a directory, relative to the current directory of the command window, corresponding to the package name of the service you are developing.
2. Alternatively, use workbench tools that support web services development. The example development tasks use Microsoft Visual Studio 2008, Eclipse IDE for Java EE Developers and WebSphere Application Server Community Edition.

### About this task

Existing web services need no modification to work with WebSphere transport for SOAP. The tools supplied with IBM MQ transport for SOAP deploy a web service, and run it using an IBM MQ SOAP listener. The tools also generate WSDL, .NET client stubs, and `.java` proxy classes to develop IBM MQ transport for SOAP clients.

Follow these steps to create a service, and prepare it for deployment and the generation of clients. Follow the steps in the related tasks to create a service using Eclipse or Microsoft Visual Studio 2008.

### Procedure

1. Develop the service using your normal development environment.
2. Test the service using an HTTP web services client
3. Follow these steps to prepare the deployment directory:
  - For Java
    - a. Copy the `.java` file defining the service interface into the deployment directory.
    - b. Copy any `.class` files for the service into the directory corresponding to the package name.
    - c. Check that the classpath can locate all the classes that are required: compile the service `.java` file using `javac`.
  - For .NET



- a. Copy the .asmx file defining the service into the deployment directory, and
- b. If you are using the code-behind model, copy any .dll files into a *deployment directory\bin* directory.

## Developing a .NET 1 or 2 service for IBM MQ transport for SOAP using Microsoft Visual Studio 2012:

Develop the SampleStockQuote web service for .NET 1 or .NET 2 using Microsoft Visual Studio 2012.

### About this task

Create the StockQuote service with a code-behind implementation using Microsoft Visual Studio 2012.

### Procedure

1. Create a template for the service, and check that it runs on HTTP.
  - a. Start Visual Studio 2012 > **File** > **New** > **Project...** Select **C# Project Type**, **.NET Framework 2**, and **ASP.NET Web Service Application**. Type the **Name:** and **Solution Name:** StockQuoteDotNet > **OK**
  - b. Right click **Service1.asmx** in the Solution Explorer > **Rename** > StockQuote.asmx.
  - c. Change the code fragment `public class Service1` to `public class StockQuote`.
  - d. Right click **StockQuote.asmx** in the Solution Explorer > **Open with...** > **XML Editor**. Change `Class="StockQuoteDotNet.Service1"` to `Class="StockQuoteDotNet.StockQuote"`
  - e. Change the code fragment `[WebService(Namespace = "http://tempuri.org/")]` to `[WebService(Namespace = "http://stock.samples/")]`.
  - f. Remove the line of code `[ToolboxItem(false)]`.
  - g. Check everything is correct so far: **Debug** > **Start Debugging (F5)**. Verify the output in Explorer.
2. Add the methods from the sample SQDNNonInline.asmx.cs, and test the service on HTTP.
  - a. Open `MQ_INSTALLATION_PATH \tools\soap\samples\dotnet\SQDNNonInline.asmx.cs` and replace the `HelloWorld` method with the four `Quote` methods; see Figure 156 on page 1235. `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed.
  - b. **Build** > **Rebuild** the solution > Right click one of the **Thread** in error > **Resolve** > Using **System.Threading**.
  - c. Press F5 to start debugging. The service is not conformant to WS-I Basic Profile v1.1. You have the choice of either changing the `WebMethod` annotation from `[SoapRpcMethod]` to `[SoapDocumentMethod]`, or removing the annotation `[WebServiceBinding(ConformsTo = WsiProfiles.BasicProfile1_1)]`.
  - d. Press F5 to verify the implementation using HTTP.
3. Generate WSDL, clients, and run the service using IBM MQ transport for SOAP.
  - a. Open a command window in the project directory tree, where the `StockQuote.asmx` is stored.
  - b. (Optional) Use `amqswdeployWMQService` to generate artifacts. The queue manager must be started:
 

```
amqswdeployWMQService -f StockQuote.asmx
-u "jms:/queue?initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=()
&destination=REQUESTDOTNET@QM1
&targetService=StockQuote.asmx"
StockQuote.asmx StockQuote.wsd
```

All artifacts are created in the `./generated` directory tree.
  - c. (Optional) Generate just the WSDL for calling the service using IBM MQ transport for SOAP.

```
amqswsd1 -u "jms:/queue?initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=()
&destination=REQUESTDOTNET@QM1
&targetService=StockQuote.asmx"
StockQuote.asmx StockQuote.wsd1
```

- a. Run the .NET listener. Either use `.\generated\server\startWMQListener.cmd` or type the command:

```
amqSOAPNETListener -u "jms:/queue?initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=()
&destination=REQUESTDOTNET@QM1
&targetService=StockQuote.asmx"
```

4. Test the service using a client generated from the WSDL, or using the clients generated by **amqdeployMQService**.

### Sample code

The sample .NET web service, `StockQuoteDotNet`, is installed in `MQ_INSTALLATION_PATH\tools\soap\samples\dotnet`. `MQ_INSTALLATION_PATH` is the directory where IBM MQ is installed. The web service binding of the published samples differ slightly from the binding used in the task. The task uses the defaults used in Microsoft Visual Studio 2012.

There are two examples of .NET Framework 1 and .NET Framework 2 web services. `StockQuoteDotNet.asmx`, is an inline service. `SQDNNoninline.asmx`, is a code-behind web service implemented by `SQDNNoninline.asmx.cs`.

`StockQuoteDotNet` has four methods:

1. `float getQuote(String symbol)`
2. `void getQuoteOneWay(String symbol).`
3. `int asyncQuote(int delay)`
4. `float getQuoteDOC(String symbol)`

---

```

<%@ WebService Language="C#" Class="StockQuoteDotNet" %>
using System;
using System.Web.Services;
using System.Web.Services.Protocols;
using System.Web.Services.Description;
using System.Threading;
[WebService (Namespace="http://stock.samples")]
public class StockQuoteDotNet {
 [WebMethod] [SoapRpcMethod(OneWay=true)]
 public void getQuoteOneWay(String symbol) {
 if (symbol.ToUpper().Equals("DELAY")) Thread.Sleep(5000);
 System.Console.WriteLine("getQuoteOneWay was invoked.");
 }
 [WebMethod] [SoapRpcMethod]
 public float getQuote(String symbol) {
 if (symbol.ToUpper().Equals("DELAY")) Thread.Sleep(10000);
 return 88.88F;
 }
 [WebMethod] [SoapRpcMethod]
 public int asyncQuote(int delay) {
 Thread.Sleep(delay);
 return delay;
 }
 [WebMethod]
 public float getQuoteDOC(String symbol) {
 return 77.77F;
 }
}

```

---

*Figure 156. Inline service: StockQuoteDotNet.asmx*

---

```

<%@ WebService Language="C#" Codebehind="SQDNNonInline.asmx.cs" Class="SQDNNonInline" %>

```

---

*Figure 157. Code-behind: Design SQDNNonInline.asmx*

---

```

using System;
using System.Web.Services;
using System.Web.Services.Protocols;
using System.Web.Services.Description;
using System.Threading;

[WebService(Namespace = "http://stock.samples")]
public class SQDNNonInline : System.Web.Services.Protocols.SoapHttpClientProtocol
{
 [WebMethod]
 [SoapRpcMethod(OneWay = true)]
 public void getNonInlineQuoteOneWay(String symbol)
 {
 if (symbol.ToUpper().Equals("DELAY")) Thread.Sleep(5000);
 System.Console.WriteLine("getNonInlineQuoteOneWay was invoked.");
 }

 [WebMethod]
 [SoapRpcMethod]
 public float getNonInlineQuote(String symbol)
 {
 if (symbol.ToUpper().Equals("DELAY")) Thread.Sleep(10000);
 return 88.88F;
 }

 [WebMethod]
 [SoapRpcMethod]
 public int asyncNonInlineQuote(int delay)
 {
 Thread.Sleep(delay);
 return delay;
 }

 [WebMethod]
 public float getNonInlineQuoteDOC(String symbol)
 {
 return 77.77F;
 }
}

```

---

*Figure 158. Code-behind: Implementation: SQDNNonInline.asmx.cs*

### **Developing a JAX-WS EJB web service for W3C SOAP over JMS:**

A web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a JEE application server. This task is step 2 of connecting an Axis2 web service client and a web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol.

#### **Before you begin**

Use Rational Application Developer to create the EJB web service. The web service wizard in Rational Application Developer has an option to create a web service using the W3C candidate recommendation for the SOAP over JMS binding. Rational Application Developer 7.54 is required. The exercise used Rational Application Developer included in Rational Software Architect for WebSphere Software v7.5.5.1,

The EJB is deployed to WebSphere Application Server from Rational Application Developer as part of this task.

To create the WSDL actually used in the task, you must first set up the Liberty profile. Then you can either import the WSDL from the Dynamic Web project in the Eclipse Galileo workspace, or from the running HTTP web service deployed to the Liberty profile.

WebSphere Application Server might still be running. If it is not, you can start it from the Servers view in RAD.

### About this task

In this task, you redeploy the StockQuoteAxis service from running as a JAX-RPC Axis service run by the **SimpleJavaListener** using IBM MQ transport for SOAP, to being a JAX-WS service running in WebSphere Application server using the W3C SOAP over JMS protocol.

There are two parts to migrating the service from the **SimpleJavaListener** to WebSphere Application Server:

1. Generate the web service interface from the WSDL for the service using the Top-down EJB web service wizard in Rational Application Developer.
2. Implementing the service by importing the IBM MQ SOAP sample StockQuoteAxis.java.

An alternative approach would have been to generate the service bottom up, from the StockQuoteAxis.java. However, to be sure that the interface to the migrated service is identical, the top-down approach is better, as it uses the same WSDL.

The web service is developed for the EJB container and not the Web container because the JMS support is part of the EJB container.

### Procedure

1. Start Rational Application Developer, and verify WebSphere Application Server is running.
  - a. Start Rational Application Developer in a new workspace.
  - b. Open the Java EE perspective.
  - c. Open the Servers tab, and check WebSphere Application Server is running.
    - If there is no WebSphere Application Server v7.0 in the view, right-click in the view > **New > Server**. Follow the choices in the wizard to create a WebSphere Application Server v7.0 instance.
    - If the server is present, but not started, click the arrow head to start it.
    - To verify the properties and get quick access to the server logs, right-click **WebSphere Application Server v7.0 at localhost > Properties > WebSphere Application Server**.
    - To administer the server, either use an external browser and open the URL, <http://localhost:9061/ibm/console/unsecureLogon.jsp>, or right-click **WebSphere Application Server v7.0 at localhost > Run administrative console**.
    - The default setting is to publish automatically. Many people prefer deploy updates to the server manually. Double-click **WebSphere Application Server v7.0 at localhost**, and expand the Publishing twisty in the Overview window. Click **Never publish automatically**.
    - Another default that you might want to alter is to clear the **Terminate server on workbench shutdown** checkbox in the Overview window.
2. Create the JEE projects You must create an Enterprise Application Project (EAR) and an Enterprise Java Bean (EJB) Project.
  - a. **File > New > Enterprise Application Project**. Name the project W3CJMSEAR > **Finish**.

The defaults must identify WebSphere Application Server v7.0 as the target runtime, and EAR version 5.0. The default configuration must be selected.
  - b. **File > New > EJB Project**. Name the project W3CJMSEJB. Select W3CEARJMS as the **EAR Project Name > Next**.

The default EJB module version is 3.0 and the default configuration is used again.
  - c. Clear the **Create an EJB Client JAR module** checkbox > **Finish**.
3. Generate and deploy the EJB web service from the StockQuoteAxis WSDL.
  - a. **Run > Launch the Web Services Explorer**.

- b. Select the WSDL page using the icons in the Web Services Explorer window > click **WSDL main** in the Navigator.
- c. In the Actions window, type in or browse to the WSDL URL for StockQuoteAxis.wsd1.  
If you have Liberty running with StockQuoteAxis deployed as an HTTP service, the URL is:  
`http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl`  
  
If you have the WSDL in the file system, the URL might be:  
`File:\Dirpath\StockQuoteAxis\WebContent\wsdl\StockQuoteAxis.wsd1`
- d. Click the line containing the imported URL in the Navigator tree.  
It is the line immediately following **WSDL Main**, if it is the first WSDL you have imported into Web Services Explorer.
- e. In the Actions window, click **Launch Web Service Wizard > Web Service Skeleton > Go**
- f. In the Web Service wizard, select **Top down EJB Web Service**  
Select or verify the Configuration using information from Table 183 Check **Overwrite files without warning > Next**.

Table 183. Top down EJB web service configuration

Field	Value
Server	WebSphere Application Server v7.0
Web service run time	IBM WebSphere JAX-WS
Service project	W3CJMSEJB
Service EAR project	W3CJMSEAR
Configuration:	No client generation

- g. On the page subtitled, Specify options for creating a WebSphere JAX-WS EJB Top Down Web Service, check the **Switch to JMS binding** box. Also check **Enable Wrapper Style**, **Copy WSDL to project**, and **Generate Web Service Deployment Descriptor > Next**.
- h. On the page titled, WebSphere JAX-WS JMS Binding Configuration, check **Use SOAP/JMS interoperability protocol** and provide values from Table 184, leaving other fields blank > **Next**.

Table 184. WebSphere JAX-WS JMS Binding Configuration

Field	Value
JMS destination	queue
Destination JNDI name:	requestaxis
JMS connection factory	qm1
Reply to Name	W3CJMSEAR
Configuration:	replyaxis

- a. On the page titled, WebSphere JAX-WS Router Project Configuration, type qm1as in the **ActivationSpec JNDI name** field > **Next**.  
RAD takes about 30 seconds to a minute to generate and deploy the project.
  - b. Ignore the options in the Web Service Publication page > **Finish**.
4. Check the generated WSDL.  
You asked for the service-specific WSDL to be generated and saved in the project.
- a. In the Enterprise Explorer navigator, open the folder **W3CJMSEJB > ejbmodule > META-INF > wsdl**. Double-click StockQuoteAxis.wsd1 to open it in the WSDL editor.  
Inspect the binding; you see the JMS url:  
`jms:jndi:requestaxis?jndiConnectionFactoryName=qm1&targetService=StockQuoteAxis`

5. Optional step: Bind the EJB to SOAP over HTTP using JAX-WS.  
Providing two bindings to the EJB gives clients the choice of SOAP bindings to call the web service. It also provides clients with the means to query the Web server to obtain its WSDL, using HTTP.  
The steps to bind an EJB to SOAP over HTTP are not included as part of the task.
6. Implement and redeploy StockQuoteAxis using the sample StockQuoteAxis.java
  - a. In the Enterprise Explorer navigator, open the folder **W3CJMSEJB** > **Services** Double-click StockQuoteAxisService to open the implementation class in a Java editor.
  - b. Open the StockQuoteAxis.java sample program in the *WebSphere MQ Installation directory\tools\soap\samples\java\server* folder > Select all the methods, but not the class name > **Copy**.
  - c. In StockQuoteAxisSoapBindingImpl.java select all the methods, but not the class name, and paste in the methods from StockQuoteAxis.java.
  - d. Add a print statement to output to the WebSphere Application Server console when the service is called. Change the getQuote(String symbol) method:
 

```
public float getQuote(String symbol) {
 System.out.println("StockQuoteAxisSoapBindingImpl called with symbol: "
 + symbol);
 return ((float) 55.25);
}
```
  - e. Correct the imports: **Source** > **Organize imports** > **Save**.
  - f. Correct the three errors due to the implementation not matching the interface.  
The errors are due to three of the methods in StockQuoteAxis.java throwing exceptions, and the WSDL for the service not containing any fault messages. The problem is diagnosed as being a mismatch between the method signatures and the method web service annotations.  
Either annotate the methods with @WebFault and regenerate the WSDL, or keep the interface unchanged, and remove the exceptions.  
To keep the interface the same, remove the three throws exception from the method signatures > **Save**.

### What to do next

“Deploying to an Axis2 client using W3C SOAP over JMS” on page 1281

## Developing IBM MQ web service clients for IBM MQ transport for SOAP

Use your normal development environment to develop web service clients for use with the IBM MQ transport for SOAP.

### Before you begin

Create the service. You can use one of the examples in “Developing web services for IBM MQ transport for SOAP” on page 1232.

Make choices about how you are going to develop, deploy and use the clients, and where to go to get the WSDL for client generation.

### Decide on your approach to developing clients and services for IBM MQ transport for SOAP.

There are two approaches.

1. Use standard development tools, develop an HTTP service and client, and then use the URL for WebSphere MQ transport for SOAP.
2. Use the tools and samples supplied with IBM MQ transport for SOAP.

If you take the HTTP route, you can run the service on an HTTP server and also run it using IBM MQ transport for SOAP. To run it using IBM MQ transport for SOAP, configure the appropriate IBM MQ listener for SOAP and set up the paths and deployment descriptors to run the service.

The tools provided by IBM MQ transport for SOAP do the configuration for you. Alternatively, you can configure the environment to run the listeners.

The tools supplied with IBM MQ transport for SOAP are useful in getting started and learning how to deploy the transport. For production work, there are benefits in using standard tools, and deploying the same service accessible to different SOAP transports.

### Decide on the type of client to develop

You must decide what type of web service client to develop. The choice depends on whether you know the service interface and the address of the service.

If the interface is known, use Axis or .NET tools to generate proxy client classes from the service interface. The proxy client classes make it easier to write a client to call the service. If the location of the service is known when you develop the client, then use the static proxy interface. If the location of the service changes, for example if the service is redeployed onto a production server, then use the dynamic proxy interface.

If the service interface is not known at the time you develop a client, on Axis, you can create a Dynamic Invocation Interface (DII) client for Axis 1.4. A DII client uses a generic interface to call any service. To pass parameters to a particular service correctly, you need to build the specific service interface programmatically. Build the interface programmatically in the client, or by loading the WSDL for the service into the client. On Axis2, you can create a Dispatch client. The Dispatch client uses a document model to describe the client request, whereas a DII client uses a call model. Both work on building the request dynamically.

### Obtain the WSDL for the service

Except for the case of the service interface being built programmatically, you must first obtain the service WSDL in order to create a web service client. The service WSDL is obtainable from three different sources:

1. Directly from the web service implementation using a tool such as **java2wsdl** (Axis) or **disco** (.NET).
2. By querying the web service using the URL: *Web service http url ?wsdl*.
3. From a file, either on a file system, or from a registry such as UDDI or WebSphere Service Registry and Repository.

**Note:** If the service is not accessible using HTTP, then the WSDL query does not work. The service itself might only be available using the IBM MQ transport for SOAP.

The WSDL generated by **amqwdeployWmqService** is not the same as WSDL generated using **java2wsdl** or **disco**. The generated WSDL is also different to any WSDL you might have started with to create the service "Top Down". On Axis, the *server-config.wsdd* deployment descriptor maps the SOAP message produced by a client to an operation and service. **amqwdeployWmqService** generates a different deployment descriptor from Eclipse.

The WSDL you use to build clients depends on how the service is deployed:

#### Deployed using **amqwdeployWmqService**

Use the WSDL generated by **amqwdeployWmqService**. Specify the `-w` flag and select `rpcLiteral` WSDL. For compatibility, you can select `rpcEncoded` WSDL. `rpcEncoded` WSDL works only with .NET and Axis 1.4 clients.

#### Manual deployment using **SimpleJavaListener**

Use one of the following WSDL files:

1. WSDL used to define the service, or stored in a repository.
2. WSDL generated from the service by **java2wsdl**.
3. WSDL queried using the URL *Web service http url ?wsdl*, if available from an HTTP server. You might run a tool such as Web Services Explorer to import the service definition directly into Eclipse.



You might need to change the URI for the service. Change it from the address of the HTTP service, to the URI for the IBM MQ transport for SOAP.

#### **Manual deployment using amqSOAPNETListener.**

Use one of the following WSDL files:

1. WSDL used to define the service, or stored in a repository.
2. WSDL obtained from the .NET service class (.asmx). using **disco**.
3. WSDL queried using the URL *web services http url ?wsdl*, if available. You might run a tool such as Web Services Explorer to import the service definition directly into Eclipse.
4. WSDL obtained by running **amqswsdl** against the .NET service class (.asmx).

You might need to change the URI for the service. Change it from the address of the HTTP service, to the URI for the IBM MQ transport for SOAP.

#### **Deployed to Windows Communication Foundation**

Obtain the service WSDL by using the URL *web service http url ?wsdl*. The service must be defined with the serviceMetadata behavior configuration as part of the service definition.

#### **Deployment to a different server platform.**

Follow the guidance provided with the platform about how to obtain the correct service WSDL.

### **About this task**

Develop clients using standard development tools. The following tasks illustrate how to build clients for .NET 1 and 2, Axis 1.4 (JAX-RPC) and Axis2 (JAX-WS). For Windows Communication Foundation, see the related tasks links.

#### **Developing a JAX-RPC client for WebSphere transport for SOAP using Eclipse:**

Develop an Axis 1.4 web service client to run using the IBM MQ transport for SOAP.

#### **Before you begin**

You must have the service available. You need to have an application server running in eclipse that supports Axis 1.4 web services. In this task we use the freely available Liberty profile. You might also use Tomcat 6, which is a smaller open source application server.

#### **About this task**

The task shows the development of three types of client for the sample StockQuoteAxis service using Eclipse running on Windows. The clients are a static and dynamic client developed using the client proxy, and a DII client.

Two alternative approaches to generating the client proxies from WSDL are illustrated:

1. Generating client proxies using **amqdeployWMQService**.
2. Importing WSDL into Eclipse, and using the web service wizard to generate the client proxies.

#### **Procedure**

1. Start the Eclipse IDE for Java EE developers.
2. Create a Java project called StockQuoteAxisClient:
  - a. Switch to the Java perspective > **File** > **New** > **Java Project**. In the **Project name** field of the Create a Java Project page type, StockQuoteAxisEclipseClient. Make sure that the execution environment is either **J2SE1-1.4** or **J2SE-1.5** > **Next**.

- b. On the Java Settings page, select the **Libraries** tab > **Add External JARs...**
  - c. Browse to `MQ_INSTALLATION_PATH/java/lib` and select all the `.jar` files > **Open**.  
`MQ_INSTALLATION_PATH` is the directory where IBM MQ is installed.
  - d. Browse to `MQ_INSTALLATION_PATH/java/lib/soap` and select all the `.jar` files > **Open**. You must have installed `axis.jar` from the IBM MQ installation media into this directory.  
`MQ_INSTALLATION_PATH` is the directory where IBM MQ is installed.
  - e. The **Library** tab now references all the `.jar` files needed to build the client > **Finish**.
3. Follow one of these two approaches to create proxies in Eclipse for the sample StockQuoteAxis web service:

- Generate the client proxies using **amqwdeployWmqService**.
  - a. Create a queue manager. For the task create QM1 as the default queue manager.
  - b. Create a working directory, `samples`. Copy the `StockQuoteAxis.java` sample program into `samples/soap/server`.
  - c. Modify `amqwsetcp.cmd` in `MQ_INSTALLATION_PATH/bin` to include the current directory in the classpath. `MQ_INSTALLATION_PATH` is the directory where IBM MQ is installed.
  - d. Open a command window in `samples` and run the modified **amqwsetcp** command
  - e. Create WSDL for the StockQuoteAxis service by running the command,
 

```
amqwdeployWmqService -f soap/server/StockQuoteAxis.java -c genAxisWsd1
-u "jms:/queue?destination=REQUESTAXIS
&initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=(connectQueueManager(QM1)binding(auto))"
```

**Remember:** Use `"/`, rather than `."` or `"\"` when using Java commands.

**Tip:** Rather than import the generated proxies into Eclipse, you can import the generated WSDL from `.samples/generated`. The resulting proxies differ in two ways:

- 1) The package names are different - which you can refactor.
  - 2) The Eclipse generated proxies include an additional helper class, `StockQuoteAxisProxy.java`
- f. Create the client proxies for the StockQuoteAxis service by running the command:
 

```
amqwdeployWmqService -f soap/server/StockQuoteAxis.java -c genProxiestoAxis
-u "jms:/queue?destination=REQUESTAXIS
&initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=(connectQueueManager(QM1)binding(auto))"
```
  - g. Import the client proxies into `StockQuoteAxisClient`:
    - 1) Right click `StockQuoteAxisClient\src` > Select **File System** > **Next** > **Browse...** > find the folder `.\samples\generated\client\remote\soap\server` > **OK**.
    - 2) Check **server** in the Import page > **Finish**.
  - h. Refactor the package name to `soap.server`.
    - 1) Right click the package containing the client proxies > **Refactor** > **Rename**. Type the **New name:** `soap.server` > leave the selected defaults for the other choices > **OK**. All the errors are fixed.

- Generate the client proxies using Eclipse.  
You have a choice of ways of obtaining the WSDL for the service. In this example, the service has been deployed to the Liberty profile and you obtain the WSDL from the Web server.
  - a. In Eclipse, switch to the Web perspective, and check that the Liberty profile is running and StockQuoteAxis is deployed and synchronized.
  - b. Import the WSDL into Web Services Explorer:
    - 1) Click the **Web Services Explorer** icon in the action-bar, or click **Run** > **Launch the Web Services Explorer**.
    - 2) Click the WSDL page icon in the Web Services Explorer to switch to the WSDL page.

- 3) Click **WSDL Main** in the Navigator window of the Web Services Explorer.
  - 4) Type in the URL of the web service, followed by ?WSDL. The URL for StockQuoteAxis, deployed in the Liberty profile, is:  
`http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl`
- c. Generate the client proxies:
- 1) In the Web Services Explorer navigator, click `http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl`.
  - 2) In the Actions window, click **Launch Web Service Wizard** > leave **Web Service Client** selected > **Go**.
  - 3) On the first page of the wizard, click the **Client** project link in the configuration > Select the **StockQuoteAxisClient** client project > **OK**.

**Tip:** The wizard window might lose focus. You need to bring it back into focus manually.

- 4) The web service runtime must be Apache Axis to generate a JAX-RPC client.
  - 5) Click **Finish**.
  - 6) Change the static URL of the service to point to the IBM MQ transport for SOAP address for the StockQuoteAxis service. You might choose to skip this step, until you have tested the client with an HTTP server.
- a) Open `StockQuoteAxisServiceLocator.java` and find the declaration for `StockQuoteAxis_address`.
  - b) Change the URL to

```
"jms:/queue?destination=REQUESTAXIS
&initialContextFactory=com.ibm.mq.jms.NoJndi
&connectionFactory=(connectQueueManager(QM1)binding(auto))"
```

**Tip:** Eclipse automatically transforms & to &amp;, and the reverse, when you copy and paste strings into .java code.

- d. Create three Java client classes, each with a main method:
- 1) Create package. Right click **StockQuoteAxisClient/src** > **New Package**. Name it `soap.client` > **Finish**.
  - 2) Select `soap.client` > **New** > **Class**. Name the class `SQASStaticClient` > Check **public static void main(string [] args)** > **Finish**
  - 3) Repeat the procedure to create `SQADynamicClient.java` and `SQADIIClient.java`
- e. Write the client code.

Figure 162 on page 1247 through Figure 166 on page 1249 provide examples of the three styles of client code. The examples use an HTTP URL to test the client using the StockQuoteAxis service deployed to an HTTP server. To run the clients against the StockQuoteAxis service deployed using IBM MQ transport for SOAP, change the URL to:

```
"jms:/queue?destination=REQUESTAXIS
connectionFactory=(connectQueueManager(QM1)binding(auto))
initialContextFactory=com.ibm.mq.jms.NoJndi
targetService=soap.server.StockQuoteAxis.java
replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE"
```

- Figure 162 on page 1247 and Figure 164 on page 1248 use the proxy generated by Eclipse, which has the extra `StockQuoteAxisproxy` helper class that makes coding a little easier.
- Figure 163 on page 1248 and Figure 165 on page 1248 use the proxy generated by **amqwdeployWMQService**.
- Figure 166 on page 1249 uses no proxy classes.

Each of the clients call `com.ibm.mq.soap.Register.extension()` to link to the IBM MQ transport for SOAP. The extension is registered in the client deployment descriptor. Client deployment to Axis 1.4 is described in “Deploying a web service client to Axis 1.4 to use IBM MQ transport for SOAP” on page 1276.

- f. Run the clients by sending the SOAP request to StockQuoteAxis hosted by the WebSphere Application Server Community Edition server configured in the workspace.
  - 1) Check that the server is running, StockQuoteAxis is deployed and synchronized.
  - 2) Select or open the client you want to test > Click **Run** in the action bar. Alternatively, click the green Run icon or right click the client in the navigator > **Run As** > **Run Configurations** ... Configure the parameters you require to run the client.
- g. Run the client using the IBM MQ transport for SOAP.

The procedure uses **amqwdeployMQService** to deploy the service, and only works with the client that uses the WSDL or proxies built by **amqwdeployMQService**. To run the client using the original WSDL, or proxies built by eclipse, deploy the service with its deployment descriptor built by Eclipse. Manually start **SimpleJavaListener** using the service port binding name as the `targetServiceName`.

  - 1) Follow the instructions in “Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using **amqwdeployMQService**” on page 1271 to deploy the service to the IBM MQ Simple Java SOAP listener. The service deployment works only for the client using the WSDL or client proxies built by **amqwdeployMQService**.
  - 2) In a command window, run **amqwclientconfig** to create the client deployment descriptor file, `client-deploy.wsdd`.
  - 3) Import `client-deploy.wsdd` into the root of the Java project you want to test using IBM MQ transport for SOAP.
    - a) Right click the Java project **StockQuoteAxisEclipseClient** > **Import** > **File system** > **Next** > **Browse...**
    - b) Browse to the directory containing `client-deploy.wsdd` > **Open** > Select the directory in the Import wizard page > check `client-deploy.wsdd`.
    - c) Verify **Into folder:** has `StockQuoteAxisEclipseClient` entered > **Finish**.
  - 4) Confirm that the working directory for running a Java application in this project is the `StockQuoteAxisEclipseClient` directory:

Right click the Java project **StockQuoteAxisEclipseClient** > **Run as....** > **Run Configurations...** > Select the (x)= **Arguments** tab > Verify that in Working Directory the **Default** radio button is checked, and the path is `StockQuoteAxisEclipseClient`. Alternatively make one of the following choices to select a different location or file containing the client configuration:

    - Check **Other:** > type a directory path of your choice.
    - In the VM arguments window, type `-Daxis.ClientConfigFile= full path to client deployment descriptor file`
  - 5) Make sure the URL is configured to point to the service deployed using IBM MQ transport for SOAP. Run the client as described in step ii.

**Tip:** Typically, you might encounter one of these errors:

- 1) Exception: No client transport named 'jms' found! .
- 2) A JMS connection error.
- 3) Exception: The AXIS engine could not find a target service to invoke! targetService is soap.server.StockQuoteAxis.java
- 4) Exception: java.lang.InstantiationException: soap.server.StockQuoteAxis

Explanations:

- 1) client-config.wsdd is not found, or does include the line <transport name="jms" pivot="java:com.ibm.mq.soap.transport.jms.WMQSender"/> in client-config.wsdd.
- 2) Possibly a build path problem - not including the .jar files in *MQ\_INSTALLATION\_PATH*/java/lib. *MQ\_INSTALLATION\_PATH* is the directory where IBM MQ is installed.
- 3) Service deployment problem, either with server-config.wsdd, or with parameters passed to **SimpleSoapListener**.
- 4) Mismatch between deployment descriptor and the implementation of the service.

If you are having difficulty running the client in Eclipse, try using a command window:

- 1) Switch to the StockQuoteAxisEclipseClient\bin directory in the workspace directory tree.
- 2) Run **amqwsetcp** and **amqwclientconfig**
- 3) Run `java soap/client/SQASStaticClient`.

### Sample JAX-RPC web service clients

The sample Java web service clients shipped with IBM MQ are installed in *MQ\_INSTALLATION\_PATH* \tools\soap\samples\java\clients. *MQ\_INSTALLATION\_PATH* is the directory where IBM MQ is installed.

#### SQAxis2Axis.java

SQAxis2Axis.java, Figure 159, is a dynamic proxy client to invoke the StockQuoteAxis service. You can override the URL of the service, which is compiled into the dynamic proxy, by providing a URL on the command line.

#### SQAxis2DotNet.java

SQAxis2DotNet.java, Figure 160 on page 1246, is a dynamic proxy client to invoke the StockQuoteDotNet service. You can override the URL of the service, which is compiled into the dynamic proxy, by providing a URL on the command line.

#### Wsd1Client.java

Wsd1Client.java, Figure 161 on page 1247, is a dynamic invocation client to invoke either the StockQuoteDotNet or StockQuoteAxis service. The client invokes the StockQuoteAxis service by default. Add the command-line option `-D` invoke the StockQuoteDotNet service and `-w` to provide a different port to the one in `.\generated\StockQuoteDotNet_Wmq.wsd1`

```
package soap.clients;
import java.net.URL;
import soap.server.*;
public class SQAxis2Axis {
 public static void main(String[] args) {
 com.ibm.mq.soap.Register.extension();
 try {
 StockQuoteAxisService locator = new StockQuoteAxisServiceLocator();
 StockQuoteAxis service = null;
 if (args.length == 0)
 service = locator.getSoapServerStockQuoteAxis_Wmq();
 else
 service = locator.getSoapServerStockQuoteAxis_Wmq(
 new java.net.URL(args[0]));
 System.out.println("Response: " + service.getQuote("XXX"));
 } catch (Exception e) {
 System.out.println("\n>>> EXCEPTION WHILE RUNNING ProxyClient DEMO <<<\n");
 e.printStackTrace();
 System.exit(2);
 }
 }
}
```

Figure 159. SQAxis2Axis.java

```

public class SQAxis2DotNet {
public static void main(String[] args) {
 com.ibm.mq.soap.Register.extension();
 try {
 StockQuoteDotNet locator = new StockQuoteDotNetLocator();
 StockQuoteDotNetSoap_PortType service = null;
 if (args.length == 0)
 service = locator.getStockQuoteDotNetSoap();
 else
 service = locator.getStockQuoteDotNetSoap(new java.net.URL(
 args[0]));
 System.out.println("Response: " + service.getQuoteDOC("XXX"));
 } catch (Exception e) {
 System.out.println("\n>>> EXCEPTION WHILE RUNNING ProxyClient DEMO <<<\n");
 e.printStackTrace();
 System.exit(2);
 }
}
}

```

*Figure 160. SQAxis2DotNet.java*

```

package soap.clients;
import com.ibm.mq.soap.*;
import org.apache.axis.utils.Options;
import java.net.URL;
import javax.xml.rpc.Call;
import javax.xml.rpc.Service;
import javax.xml.rpc.ServiceFactory;
import javax.xml.namespace.QName;
public class WsdIClient {
public static void main(String[] args) {
 String wsdlService, wsdlPort, namespace, wsdlSource, wsdlTargetURI, s;
 try {
 Register.extension();
 Options opts = new Options(args);
 if (opts.isFlagSet('D') != 0) {
 wsdlService = "StockQuoteDotNet";
 wsdlPort = "StockQuoteDotNetSoap";
 namespace = "http://stock.samples";
 wsdlSource = "file:generated/StockQuoteDotNet_Wmq.wsdl";
 } else {
 wsdlService = "StockQuoteAxisService";
 wsdlPort = "soap.server.StockQuoteAxis_Wmq";
 namespace = "soap.server.StockQuoteAxis_Wmq";
 wsdlSource = "file:generated/soap.server.StockQuoteAxis_Wmq.wsdl";
 }
 if (null != (s = (opts.isValueSet('w'))))
 wsdlPort = s;
 System.out.println("start WsdIClient demo, wsdl port " + wsdlPort
 + " resolving uri to ...");
 QName servQN = new QName(namespace, wsdlService);
 QName portQN = new QName(namespace, wsdlPort);
 Service service = ServiceFactory.newInstance().createService(
 new URL(wsdlSource), servQN);
 Call call = (Call) service.createCall(portQN, "getQuote");
 wsdlTargetURI = call.getTargetEndpointAddress().toString();
 System.out.println(" " + wsdlTargetURI + "");
 Object ret = call.invoke(new Object[] { "XXX" });
 System.out.println("Response: " + ret);
 } catch (Exception e) {
 System.out.println("\n>>> EXCEPTION WHILE RUNNING WsdIClient DEMO <<<\n");
 e.printStackTrace();
 System.exit(2);
 }
}
}
}

```

Figure 161. WsdIClient.java

The example clients used in this task:

```

package soap.client;
import soap.server.StockQuoteAxisProxy;
public class SQAStaticClient {
 public static void main(String[] args) {
 try {
 com.ibm.mq.soap.Register.extension();
 StockQuoteAxisProxy sqa = new StockQuoteAxisProxy();
 System.out.println("Static client synchronous result is:"
 + sqa.getQuote("ibm"));
 } catch (Exception e) {
 System.out.println("Exception: " + e);
 }
 }
}
}

```

Figure 162. Static client using Eclipse generated proxy

```

package soap.client;
import soap.server.StockQuoteAxis;
import soap.server.StockQuoteAxisService;
import soap.server.StockQuoteAxisServiceLocator;
public class SQAStaticClient {
 public static void main(String[] args) {
 try {
 com.ibm.mq.soap.Register.extension();
 StockQuoteAxisService locator = new StockQuoteAxisServiceLocator();
 StockQuoteAxis sqa = locator.getSoapServerStockQuoteAxis_Wmq();
 System.out.println("Static client synchronous result is: "
 + sqa.getQuote("ibm"));
 } catch (Exception e) {
 System.out.println("Exception: " + e);
 }
 }
}

```

*Figure 163. Static client using amqwdeployWMQService generated proxy*

```

package soap.client;
import soap.server.StockQuoteAxisProxy;
public class SQADynamicClient {
 public static void main(String[] args) {
 try {
 com.ibm.mq.soap.Register.extension();
 StockQuoteAxisProxy sqa = new StockQuoteAxisProxy(
 "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis");
 System.out.println("Dynamic client synchronous result is: "
 + sqa.getQuote("ibm"));
 } catch (Exception e) {
 System.out.println("Exception: " + e);
 }
 }
}

```

*Figure 164. Dynamic client using Eclipse generated proxy*

```

package soap.client;

import java.net.URL;
import soap.server.StockQuoteAxis;
import soap.server.StockQuoteAxisService;
import soap.server.StockQuoteAxisServiceLocator;
public class SQADynamicClient {
 public static void main(String[] args) {
 try {
 com.ibm.mq.soap.Register.extension();
 URL sqaURL = new URL(
 "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis");
 StockQuoteAxisService locator = new StockQuoteAxisServiceLocator();
 StockQuoteAxis sqa = locator.getSoapServerStockQuoteAxis_Wmq(sqaURL);
 System.out.println("Dynamic client synchronous result is: "
 + sqa.getQuote("ibm"));
 } catch (Exception e) {
 System.out.println("Exception: " + e);
 }
 }
}

```

*Figure 165. Dynamic client using amqwdeployWMQService generated proxy*



```

package soap.client;
import java.net.URL;
import javax.xml.namespace.QName;
import javax.xml.rpc.Call;
import javax.xml.rpc.Service;
import javax.xml.rpc.ServiceFactory;
public class SQADIIClient {
 public static void main(String[] args) {
 try {
 com.ibm.mq.soap.Register.extension();
 URL wsdl = new URL(
 "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl");
 Service SQAService = (ServiceFactory.newInstance()).createService(wsdl,
 new QName("http://server.soap", "StockQuoteAxisService"));
 Call SQACall = SQAService.createCall(new QName("http://server.soap",
 "StockQuoteAxis"), "getQuote");
 System.out.println("DII client synchronous result is "
 + SQACall.invoke(new Object[] { "ibm" }));
 } catch (Exception e) {
 System.out.println("Exception: " + e);
 }
 }
}

```

Figure 166. DII client (No proxy)

### Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse:

Develop an Axis2 web service client to run using the IBM MQ transport for SOAP. The sample Axis2 clients provided with IBM MQ transport for SOAP are listed, and the **wsimport** command used to generate proxies.

#### Before you begin

Obtain the Axis2 libraries, and configure a development and test environment to run the client.

**Note:** The naming of versions and releases used by Axis causes confusion. Typically, Axis 1.4 refers to the JAX-RPC implementation, and Axis2 to the JAX-WS implementation.

Axis 1.4 is a version level. If you search for Axis 1.4 on the internet, you are taken to <http://ws.apache.org/axis/>. The page contains a list of preceding versions of Axis (1.2, 1.3) and the April 22, 2006, final release of Axis 1.4. There are later releases of Axis 1.4, that fix bugs, but they are all known as Axis 1.4. It is one of these bug fix releases that is shipped with IBM MQ. For Axis 1.4, use the version of `axis.jar` that is shipped with IBM MQ rather than the one obtainable from <http://ws.apache.org/axis/>.

The Axis website also refers to Axis 1.1 to refer to all the versions of what is more typically called Axis 1.4. Axis 1.2 is used to refer to what is typically called Axis2.

Axis 1.5 is not a later release of Axis 1.4, it is an Axis2 release. If you search for Axis 1.5 you are directed to <http://ws.apache.org/axis2/>. contains a list of release versions of Axis2, labeled 0.9 to 1.5.1 (and including, confusingly version 1.4). The release version of Axis2 to use with IBM MQ transport for SOAP is 1.4.1. Download Axis2 1.4.1 from [http://ws.apache.org/axis2/download/1\\_4\\_1/download.cgi](http://ws.apache.org/axis2/download/1_4_1/download.cgi).

You can choose to generate proxies for the web service clients for IBM MQ transport for SOAP using either **wsimport** or the tooling provided with an IDE. Eclipse IDE for Java EE Developer 3.5 SR1 uses **wsdl2java**. **wsimport** is supplied with Java 6. You can use Java 5 to run client proxies generated either with **wsimport** or **wsdl2java**.

The sample web service Axis2 clients provided with IBM MQ transport for SOAP were developed using **wsimport** ; see “Sample Axis2 clients” on page 1256.

The task that follows demonstrates how to generate and use the proxies produced by the web services wizard that is packaged with Eclipse IDE for Java EE Developers. The sample clients show how to use the proxies produced by **wsimport**.

To use the web services wizard, you must add an application server that supports Axis2 to the workbench. The steps show how to configure the Liberty profile to support Axis2 using the workbench.

1. Configure the application server used in Eclipse IDE for Java EE Developers to support Axis2. In this example, configure the Liberty profile.
  - a. Open the workspace preferences to configure the server: Open **Window > Preferences**.
  - b. Check the installed JRE is Java50: Click **Installed JREs**.
  - c. Add the Liberty profile as the server:
  - d. Add Axis2: Click **Web Services > Axis2 Preferences**. On the Axis2 Runtime tab > **Browse...** Open the directory containing the many Axis2 jar files > **Apply**.
  - e. Associate Liberty with Axis2: Click **Web Services > Server and Runtime**. Under **Server** select **IBM Liberty Server**, and under **Web service runtime**, select **Apache Axis2 > Apply > OK**
  - f. Start the server: Open the Web perspective and open the Servers view. Right click in the Servers view > **New > Server**. **IBM Liberty Server** is selected and configured > **Finish**. Start the server.
2. Check that you have deployed the StockQuoteAxis service to Liberty to run the web service wizard.
3. To test the service with the IBM MQ transport for SOAP service, deploy the service to an IBM MQ transport for SOAP listener for Axis 1.4; see the Liberty profile.

### About this task

The Eclipse IDE for Java EE Developers uses Java50 and the web services wizard to generate the proxy classes for the service. The proxy classes are different to the classes created by the **wsimport** tool provided with Java 6. An alternative approach is to generate the proxy classes using **wsimport** and import the packages it creates into your Eclipse Java EE IDE for Web Developers.

The web services wizard in the Eclipse IDE for Java EE Developers builds a web service client in a Web project. You can run the client as a simple Java application; it does not require an application server. You can also transfer the code to a Java project, and configure the build path to include the Axis2 JAR files.

### Procedure

1. Create a Web project in a new Enterprise project:
  - a. With nothing selected in the Project Explorer > Right-click the white space > **New > Enterprise Application Project** > Name it StockQuoteAxis2EAR > **Finish**. Reply No to the window giving you the option of opening the Java EE perspective. The defaults are set to use Liberty.
  - b. Right-click StockQuoteAxis2EAR > **New > Dynamic Web Project**. Name the project StockQuoteAxis2WebClient > Check the EAR membership box to add the project to **StockQuoteAxis2EAR**. Liberty is selected as the Target runtime.
  - c. In the Configuration section of the New Dynamic Web Project page > **Modify...** > Check the Axis2 web services project facet. **Dynamic Web Module 2.5, Java 6.0**, and **Liberty** are already checked. > **OK > Finish**. Reply No to the window giving you the option of opening the Java EE perspective.
2. Import WSDL for the service into the workspace and generate the client proxy:

In this example, the WSDL document contains the HTTP service binding and becomes the target for the static Web client proxy. You can modify the URL in the web service binding to point to the IBM MQ transport for SOAP URL before generating the client proxy. The static Web client proxy is then the service that is deployed to IBM MQ transport for SOAP.

- a. Launch the Web Services Explorer: either use the icon in the action bar, or **Run > Launch the Web Services Explorer**.
  - b. Select the WSDL explorer by clicking the WSDL icon in the Web Services Explorer window > Click **WSDL Main** in the Navigator window > Type the URL of the StockQuoteAxis WSDL file > **Go**. In this example, obtain the WSDL directly from the HTTP service: `http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl`
  - c. In the Navigator, click the line with the URL of the web service. In the Actions window, click **Import WSDL to Workbench** > Select a **StockQuoteAxis2WebClient** as the **Workbench project** > Type the **WSDL file name**, `StockQuoteAxisHTTP.wsdl` > **Go**.
  - d. Right-click **StockQuoteAxisHTTP.wsdl** > **Web Services** > **Generate Client**. Check the configuration information about the web services page of the wizard is as follows: Server: IBM Liberty Server, web service runtime: Apache Axis2, Client project: StockQuoteAxis2WebClient, Client EAR project: StockQuoteAxisEAR. To correct the configuration, click the lines that are wrong.
  - e. Click **Next** > verify the code generation settings > **Finish**. Notice that a new package, `soap.server`, is created and it contains the proxies you require.
3. Configure the project to run IBM MQ transport for SOAP as the JMS transport. IBM MQ transport for SOAP provides a `transportSender`, but no `transportReceiver`. In other words, IBM MQ transport for SOAP supports Axis2 clients. Currently it does not support Axis2 services.
    - a. In the **StockQuoteAxis2WebClient** project, right-click `WebContent\WEB-INF\conf\axis2.xml` > **Open with...** > **XML editor**.
    - b. Search for the last `transportSender` (towards the end of the file) and find the commented out JMS `transportSender` > Right-click the line > **Add before...** > **transportSender**.
    - c. Right-click **transportSender** > **Add Attribute** > **Name** > Right-click **transportSender** > **Add Attribute** > **Class**.
    - d. Right-click **Name** > **Edit Attribute** > Type the **Value**: `jms`
    - e. Right-click **Class** > **Edit Attribute** > Type the **Value**: `com.ibm.mq.axis2.transport.jms.WMQJMSTransportSender`.> **Save**.
    - f. Add `com.ibm.mq.axis2.transport.jms.WMQJMSTransportSender` to the build path: Right-click **StockQuoteAxis2WebClient** > **Build Path** > **Configure Build Path...** > Click the Libraries tab > **Add External JARs...** Select all the JARs in `MQ_INSTALLATION_PATH\java\lib` > **OK**. `MQ_INSTALLATION_PATH` is the directory where IBM MQ is installed.
  4. Create a synchronous static client, test it using HTTP, then convert the proxy to run the static client using IBM MQ transport for SOAP.
    - a. Right-click **Java Resources: src** > **New** > **Package** > Name the package `soap.client` > **Finish**
    - b. Right-click **soap.client** > **New** > **Class** > Name the class `SQA2StaticClient` > **Finish**.
    - c. Replace the class with the code in Figure 167 on page 1252 > **Save**.

---

```

package soap.client;
import soap.server.StockQuoteAxisServiceStub;
import soap.server.StockQuoteAxisServiceStub.GetQuote;
public class SQA2StaticClient {
 public static void main(String[] args) {
 try {
 StockQuoteAxisServiceStub stub = new StockQuoteAxisServiceStub();
 GetQuote request = new GetQuote();
 request.setSymbol("ibm");
 System.out.println("Response is: "
 + (stub.getQuote(request)).getGetQuoteReturn());
 } catch (Exception e) {
 System.out.println("Exception: " + e.getMessage());
 e.printStackTrace();
 }
 }
}

```

---

Figure 167. SQA2StaticClient.java

5. Test the client with the StockQuoteAxis service deployed to Liberty, and with IBM MQ transport for SOAP.
  - a. In the Project Explorer, right-click **SQA2StaticClient** > **Run as... > Java Application**. The result, Response is 55.25 , appears in the Console view. You can also select the Liberty console window in the Console view, and see the output on the Liberty server, StockQuoteAxis called with parameter: ibm .
  - b. The proxy was built with the service address, `http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis`, and so the static client calls the service running on HTTP. You can change the static client to call the service using IBM MQ transport for SOAP. The following instructions change the service address in `StockQuoteAxisServiceStub.java` without rebuilding the proxy, and configure the `SQA2StaticClient` runtime parameters to load `axis2.xml`. You configure `axis2.xml` configures Axis2 to use IBM MQ transport for SOAP.
  - c. Open `StockQuoteAxisServiceStub.java` > Replace the two occurrences of `http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis` with,
 

```

jms:/queue?destination=REQUESTAXIS@QM1
&connectionFactory=()
&initialContextFactory=com.ibm.mq.jms.NoJndi
&targetService=StockQuoteAxis
&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE

```
  - d. If you run `SQA2StaticClient` now, it throws an exception because it has not found a `transportSender` configured for JMS The exception is:
 

```

Exception: null java.lang.NullPointerException
soap.server.StockQuoteAxisServiceStub.getQuote(StockQuoteAxisServiceStub.java:547)
at soap.client.SQA2StaticClient.main(SQA2StaticClient.java:11)

```
  - e. In the Project Explorer, right-click **SQA2StaticClient** > **Run as... > Run Configurations...** Switch to the (x)= **Arguments** tab, and in the **VM arguments** input area, type the path to the `axis2.conf` file> **Apply** > **Run**. The VM argument is:
 

```

-Daxis2.xml=${workspace_loc:StockQuoteAxis2WebClient/WebContent/WEB-INF/conf}/axis2.xml.

```

 Or you can provide a standard path to the Axis2 configuration file.
  - f. Run `SQA2StaticClient` again. On this run, you are using the IBM MQ transport for SOAP. Confirm it by checking there is no new output in the Liberty console. Open the console or command window that is associated with `SimpleJavaListener`, and the output there is `StockQuoteAxis` called with parameter: ibm .
6. Create a dynamic client for HTTP and IBM MQ transport for SOAP, and test it.
  - a. Right-click **soap.client** > **New > Class > Name the class SQA2DynamicClient > Finish**.
  - b. Replace the class with the code in Figure 168 on page 1253 > **Save**.

---

```

package soap.client;
import soap.server.StockQuoteAxisServiceStub;
import soap.server.StockQuoteAxisServiceStub.GetQuote;
public class SQA2DynamicClient {
 public static void main(String[] args) {
 try {
 StockQuoteAxisServiceStub stub = new StockQuoteAxisServiceStub(
 "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis");
 GetQuote request = new GetQuote();
 request.setSymbol("ibm");
 System.out.println("HTTP Sync: "
 + (stub.getQuote(request)).getGetQuoteReturn());
 stub = new StockQuoteAxisServiceStub(
 "jms:/queue?destination=REQUESTAXIS@QM1"
 + "&connectionFactory=()&initialContextFactory=com.ibm.mq.jms.NoJndi"
 + "&targetService=StockQuoteAxis&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE");
 System.out.println("JMS sync: "
 + (stub.getQuote(request)).getGetQuoteReturn());
 } catch (Exception e) {
 System.out.println("Exception: " + e.getMessage());
 e.printStackTrace();
 }
 }
}

```

---

Figure 168. SQA2DynamicClient.java

- c. Create a Run configuration for SQA2DynamicClient.java, and add the path to axis2.xml:
    - Daxis2.xml=\${workspace\_loc:StockQuoteAxis2WebClient/WebContent/WEB-INF/conf}/axis2.xml
  - d. Run SQA2DynamicClient. Check the console output for the SQA2DynamicClient, Liberty and SimpleJavaListener.
7. Create an asynchronous client, and access the result in a callback handler, and in the main program thread.

The asynchronous client proxies created by the web service wizard for Eclipse Java EE IDE for Web Developers differ from the proxies created by **wsimport**. The **wsimport** proxies use Future, Response, and AsyncHandler generic types.

The web service wizard for Eclipse Java EE IDE for Web Developers creates a StockQuoteAxisServiceCallbackHandler abstract class. You must extend StockQuoteAxisServiceCallbackHandler and create a callback handler.

- a. Right-click **soap.client** > **New** > **Class** > Name the class SQA2CallbackHandler > **Finish**.
- b. Replace the class with the code in Figure 169 on page 1254.

---

```
package soap.client;
import soap.server.StockQuoteAxisServiceCallbackHandler;
import soap.server.StockQuoteAxisServiceStub.GetQuoteResponse;
public class SQA2CallbackHandler
 extends StockQuoteAxisServiceCallbackHandler {
 private boolean complete = false;
 SQA2CallbackHandler() {
 super();
 System.out.println("Callback constructor");
 }
 public void receiveResultgetQuote(GetQuoteResponse response) {
 System.out.println("Result in Callback " + response.getGetQuoteReturn());
 super.clientData = response;
 complete = true;
 }
 public boolean isComplete() {
 return complete;
 }
}
```

---

*Figure 169. SQA2CallbackHandler.java*

- c. Right-click **soap.client** > **New** > **Class** > Name the class **SQA2AsyncClient** > **Finish**.
- d. Replace the class with the code in Figure 170 on page 1255.

---

```

package soap.client;
import soap.server.StockQuoteAxisServiceStub;
import soap.server.StockQuoteAxisServiceStub.GetQuote;
import soap.server.StockQuoteAxisServiceStub.GetQuoteResponse;
import soap.server.StockQuoteAxisServiceCallbackHandler;
@SuppressWarnings("unused")
public class SQA2AsyncClient {
 public static void main(String[] args) {
 try {
 StockQuoteAxisServiceStub stub = new StockQuoteAxisServiceStub(
 "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis");
 GetQuote request = new GetQuote();
 request.setSymbol("ibm");
 System.out.println("HTTP Sync: "
 + (stub.getQuote(request)).getGetQuoteReturn());
 SQA2CallbackHandler callback = new SQA2CallbackHandler();
 stub.startgetQuote(request, callback);
 do {
 System.out.println("Waiting for HTTP callback");
 Thread.sleep(2000);
 } while (!callback.isComplete());
 System.out.println("HTTP poll: "
 + ((GetQuoteResponse) (callback.getClientData()))
 .getGetQuoteReturn());
 stub = new StockQuoteAxisServiceStub(
 "jms:/queue?destination=REQUESTAXIS@QM1"
 + "&connectionFactory=()&initialContextFactory=com.ibm.mq.jms.NoJndi"
 + "&targetService=StockQuoteAxis&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE");
 System.out.println("JMS Sync: "
 + (stub.getQuote(request)).getGetQuoteReturn());
 callback = new SQA2CallbackHandler();
 stub.startgetQuote(request, callback);
 while (!callback.isComplete()) {
 System.out.println("Waiting for JMS callback");
 Thread.sleep(2000);
 }
 System.out.println("JMS poll: "
 + ((GetQuoteResponse) (callback.getClientData())).getGetQuoteReturn());
 } catch (Exception e) {
 System.out.println("Exception: " + e.getMessage());
 e.printStackTrace();
 }
 }
}

```

---

*Figure 170. SQA2AsyncClient.java*

The console output is listed in Figure 171.

*Figure 171. Console output from SQA2AsyncClient.java*

```

HTTP Sync: 55.25
Callback constructor
Waiting for HTTP callback
Result in Callback 55.25
HTTP poll: 55.25
JMS Sync: 55.25
Callback constructor
Waiting for JMS callback
Result in Callback 55.25
JMS poll: 55.25

```

## Sample Axis2 clients

The sample proxies are generated using the **wsimport** tool that is packaged with Java 6. Six samples are provided:

1. `DynamicProxyClientSync.java`
2. `DynamicProxyClientAsyncPolling.java`
3. `DynamicProxyClientAsyncCallback.java`
4. `DispatchClientSync.java`
5. `DispatchClientAsyncPolling.java`
6. `DispatchClientAsyncCallback.java`

The client samples are generated for the sample `StockQuoteAxis` server. Generate the WSDL with the **amqwdpoyMQServer** command, specifying the `-w` switch to select `rpcLiteral` style. Use the following command to generate the proxies for the samples:

```
wsimport soap.server.StockQuoteAxis_Wmq.wsdl -d generated -keep -p com.ibm.mq.axis2.samples
```



---

```

package com.ibm.mq.axis2.samples;

import com.ibm.mq.axis2.samples.proxy.StockQuoteAxis;
import com.ibm.mq.axis2.samples.proxy.StockQuoteAxisService;

public class DynamicProxyClientSync {

 public static void main(String[] args) {
 try {
 System.out.println("Starting sample DynamicProxyClientSync");

 System.out.println("Creating proxy instance for service StockQuoteAxisService");
 StockQuoteAxisService stub = new StockQuoteAxisService();
 StockQuoteAxis service = stub.getSoapServerStockQuoteAxisWmq();

 System.out.println("Invoking getQuoteOneWay OneWay operation synchronously...");
 service.getQuoteOneWay("48");
 System.out.println(" > getQuoteOneWay has returned");

 System.out.println("Invoking getQuote Request Reply operation synchronously...");
 float result = service.getQuote("48");
 System.out.println(" > getQuote has returned result of " + result);

 System.out.println("End of sample");
 }
 catch (Exception fault) {
 // Identify the cause of the Axis Fault
 System.err.println(fault.toString());
 Throwable e = fault.getCause();
 for (int i = 1; e != null; i++) {
 // The toString method on an MQAxisException will cause the message, explanation and user
 // action.
 System.err.println("Exception(" + i + "): " + e.toString());

 if (e.getCause() != null) {
 e = e.getCause();
 }
 else {
 break;
 }
 } // end of for loop
 } // end of catch block
 }
}

```

---

*Figure 172. DynamicProxyClientSync.java*

---

```

package com.ibm.mq.axis2.samples;

import java.util.concurrent.CancellationException;

import javax.xml.ws.Response;

import com.ibm.mq.axis2.samples.proxy.StockQuoteAxis;
import com.ibm.mq.axis2.samples.proxy.StockQuoteAxisService;

public class DynamicProxyClientAsyncPolling {

 public static void main(String[] args) {
 try {
 System.out.println("Starting sample DynamicProxyClientAsyncPolling");

 System.out.println("Creating proxy instance for service StockQuoteAxisService");
 StockQuoteAxisService stub = new StockQuoteAxisService();
 StockQuoteAxis service = stub.getSoapServerStockQuoteAxisWmq();

 System.out
 .println("Invoking getQuoteAsync Request Reply operation asynchronously by polling...");
 Response<Float> response = service.getQuoteAsync("49");

 /** Sleep main thread until response arrives **/
 System.out.println("Waiting for response to arrive...");
 while (!response.isDone()) {
 Thread.sleep(100);
 }
 System.out.println(" > Response received");

 /** Retrieve the result **/
 try {
 Float result = response.get();
 System.out.println("> getQuoteAsync call has returned result of " + result);
 }
 catch (CancellationException ce) {
 // processing was cancelled via response.cancel()
 }

 System.out.println("End of sample");
 }
 catch (Exception fault) {
 // Identify the cause of the Axis Fault
 System.err.println(fault.toString());
 Throwable e = fault.getCause();
 for (int i = 1; e != null; i++) {
 // The toString method on an MQAxisException will cause the message, explanation and user
 // action.
 System.err.println("Exception(" + i + "): " + e.toString());

 if (e.getCause() != null) {
 e = e.getCause();
 }
 else {
 break;
 }
 } // end of for loop
 } // end of catch block
 }
}

```

---

*Figure 173. DynamicProxyClientAsyncPolling.java*

---

```

package com.ibm.mq.axis2.samples;

import java.util.concurrent.Future;

import javax.xml.ws.AsyncHandler;
import javax.xml.ws.Response;

import com.ibm.mq.axis2.samples.proxy.StockQuoteAxis;
import com.ibm.mq.axis2.samples.proxy.StockQuoteAxisService;

public class DynamicProxyClientAsyncCallback implements AsyncHandler<Float> {

 public static void main(String[] args) {
 try {
 System.out.println("Starting sample DynamicProxyClientAsyncCallback");

 System.out.println("Creating proxy instance for service StockQuoteAxisService");
 StockQuoteAxisService stub = new StockQuoteAxisService();
 StockQuoteAxis service = stub.getSoapServerStockQuoteAxisWmq();

 DynamicProxyClientAsyncCallback handler = new DynamicProxyClientAsyncCallback();

 System.out
 .println("Invoking getQuoteAsync Request Reply operation asynchronously using a callback...");
 Future<?> monitor = service.getQuoteAsync("50", handler);
 System.out.println(" > Invoke call has returned");

 /** Sleep main thread until handler has been notified **/
 System.out.println("Waiting for handler to be called...");
 while (!monitor.isDone()) {
 Thread.sleep(100);
 }

 System.out.println("End of sample");
 }
 catch (Exception fault) {
 // Identify the cause of the Axis Fault
 System.err.println(fault.toString());
 Throwable e = fault.getCause();
 for (int i = 1; e != null; i++) {
 // The toString method on an MQAxisException will cause the message, explanation and user
 // action.
 System.err.println("Exception(" + i + "): " + e.toString());

 if (e.getCause() != null) {
 e = e.getCause();
 }
 else {
 break;
 }
 } // end of for loop
 } // end of catch block
 }

 public void handleResponse(Response<Float> response) {
 try {
 Float result = response.get();
 System.out.println(" > Async Handler has received a result of " + result);
 }
 catch (Exception fault) {
 // Identify the cause of the Axis Fault
 System.err.println("Exception in handleResponse");
 System.err.println(fault.toString());
 Throwable e = fault.getCause();
 for (int i = 1; e != null; i++) {
 // The toString method on an MQAxisException will cause the message, explanation and user
 // action.
 System.err.println("Exception(" + i + "): " + e.toString());

 if (e.getCause() != null) {
 e = e.getCause();
 }
 }
 }
 }
}

```



```

package com.ibm.mq.axis2.samples;

import javax.xml.namespace.QName;
import javax.xml.soap.MessageFactory;
import javax.xml.soap.SOAPBody;
import javax.xml.soap.SOAPConstants;
import javax.xml.soap.SOAPElement;
import javax.xml.soap.SOAPEnvelope;
import javax.xml.soap.SOAPHeader;
import javax.xml.soap.SOAPMessage;
import javax.xml.soap.SOAPPart;
import javax.xml.ws.Dispatch;
import javax.xml.ws.Service;
import javax.xml.ws.soap.SOAPBinding;

public class DispatchClientSync {

 public static void main(String[] args) {
 try {
 System.out.println("Starting sample DispatchClientSync");

 String endpointUrl = "jms:/queue?destination=SOAPJ.demos@WMQSOAP.DEMO.QM&"
 + "connectionFactory=connectQueueManager(WMQSOAP.DEMO.QM)"
 + "&initialContextFactory=com.ibm.mq.jms.NoJndi&targetService=soap.server.StockQuoteAxis.java";

 QName serviceName = new QName("soap.server.StockQuoteAxis_Wmq", "StockQuoteAxisService");
 QName portName = new QName("soap.server.StockQuoteAxis_Wmq", "soap.server.StockQuoteAxis_Wmq");

 Service service = Service.create(serviceName);
 service.addPort(portName, SOAPBinding.SOAP11HTTP_BINDING, endpointUrl);

 /** Create a Dispatch instance from a service */
 System.out.println("Creating dispatch instance for service StockQuoteAxisService");
 Dispatch<SOAPMessage> dispatch = service.createDispatch(portName, SOAPMessage.class,
 Service.Mode.MESSAGE);
 System.out.println(" > Dispatch instance created.");

 /*****
 * Create OneWay SOAPMessage request.
 *****/
 MessageFactory mf = MessageFactory.newInstance(SOAPConstants.SOAP_1_1_PROTOCOL);

 System.out.println("\nCreating a OneWay SOAP Message");
 SOAPMessage request = mf.createMessage();

 /** Obtain the SOAPEnvelope and header and body elements */
 SOAPPart part = request.getSOAPPart();
 SOAPEnvelope env = part.getEnvelope();
 SOAPHeader header = env.getHeader();
 SOAPBody body = env.getBody();

 /** Construct the message payload */
 SOAPElement operation = body.addChildElement("getQuoteOneWay", "ns1",
 "soap.server.StockQuoteAxis_Wmq");
 SOAPElement value = operation.addChildElement("in0");
 value.addAttribute(new QName("http://www.w3.org/2001/XMLSchema-instance", "type"), "string");
 value.addTextNode("XXX");
 request.saveChanges();
 System.out.println(" > SOAP Message created.");

 /** Invoke the service endpoint */
 System.out.println("Invoking getQuoteOneWay OneWay operation synchronously...");
 dispatch.invokeOneWay(request);
 System.out.println(" > getQuoteOneWay call has returned");

 /*****
 * Create Request Reply SOAPMessage request.
 *****/
 mf = MessageFactory.newInstance(SOAPConstants.SOAP_1_1_PROTOCOL);

 System.out.println("\nCreating a Request Reply SOAP Message");
 request = mf.createMessage();

```



```

package com.ibm.mq.axis2.samples;

import javax.xml.namespace.QName;
import javax.xml.soap.MessageFactory;
import javax.xml.soap.SOAPBody;
import javax.xml.soap.SOAPConstants;
import javax.xml.soap.SOAPElement;
import javax.xml.soap.SOAPEnvelope;
import javax.xml.soap.SOAPHeader;
import javax.xml.soap.SOAPMessage;
import javax.xml.soap.SOAPPart;
import javax.xml.ws.Dispatch;
import javax.xml.ws.Response;
import javax.xml.ws.Service;
import javax.xml.ws.soap.SOAPBinding;

public class DispatchClientAsyncPolling {

 public static void main(String[] args) {
 try {
 System.out.println("Starting sample DispatchClientAsyncPolling");

 String endpointUrl = "jms:/queue?destination=SOAPJ.demos@WMQSOAP.DEMO.QM&"
 + "connectionFactory=connectQueueManager(WMQSOAP.DEMO.QM)"
 + "&initialContextFactory=com.ibm.mq.jms.NoJndi&targetService=soap.server.StockQuoteAxis.java";

 QName serviceName = new QName("soap.server.StockQuoteAxis_Wmq", "StockQuoteAxisService");
 QName portName = new QName("soap.server.StockQuoteAxis_Wmq", "soap.server.StockQuoteAxis_Wmq");

 Service service = Service.create(serviceName);
 service.addPort(portName, SOAPBinding.SOAP11HTTP_BINDING, endpointUrl);

 /** Create a Dispatch instance from a service. */
 System.out.println("Creating dispatch instance for service StockQuoteAxisService");
 Dispatch<SOAPMessage> dispatch = service.createDispatch(portName, SOAPMessage.class,
 Service.Mode.MESSAGE);
 System.out.println(" > Dispatch instance created.");

 /** Create SOAPMessage request. */
 MessageFactory mf = MessageFactory.newInstance(SOAPConstants.SOAP_1_1_PROTOCOL);

 System.out.println("Creating a Request Reply SOAP Message");
 SOAPMessage request = mf.createMessage();

 /** Obtain the SOAPEnvelope and header and body elements */
 SOAPPart part = request.getSOAPPart();
 SOAPEnvelope env = part.getEnvelope();
 SOAPHeader header = env.getHeader();
 SOAPBody body = env.getBody();

 /** Construct the message payload */
 SOAPElement operation = body.addChildElement("getQuote", "ns1",
 "soap.server.StockQuoteAxis_Wmq");
 SOAPElement value = operation.addChildElement("in0");
 value.addAttribute(new QName("http://www.w3.org/2001/XMLSchema-instance", "type"), "string");
 value.addTextNode("XXX");
 request.saveChanges();
 System.out.println(" > SOAP Message created.");

 /** Invoke the service endpoint */
 System.out.println("Invoking getQuote Request Reply operation asynchronously by polling...");
 Response<SOAPMessage> response = dispatch.invokeAsync(request);
 System.out.println(" > getQuote call has returned");

 /** Sleep main thread until response arrives */
 System.out.println("Waiting for response to arrive...");
 while (!response.isDone()) {
 Thread.sleep(100);
 }
 System.out.println(" > Response received");

 /** retrieve the result */

```





```

package com.ibm.mq.axis2.samples;

import java.util.concurrent.Future;

import javax.xml.namespace.QName;
import javax.xml.soap.MessageFactory;
import javax.xml.soap.SOAPBody;
import javax.xml.soap.SOAPConstants;
import javax.xml.soap.SOAPElement;
import javax.xml.soap.SOAPEnvelope;
import javax.xml.soap.SOAPHeader;
import javax.xml.soap.SOAPMessage;
import javax.xml.soap.SOAPPart;
import javax.xml.ws.AsyncHandler;
import javax.xml.ws.Dispatch;
import javax.xml.ws.Response;
import javax.xml.ws.Service;
import javax.xml.ws.soap.SOAPBinding;

public class DispatchClientAsyncCallback implements AsyncHandler<SOAPMessage> {

 public static void main(String[] args) {
 try {
 System.out.println("Starting sample DispatchClientAsyncCallback");

 String endpointUrl = "jms:/queue?destination=SOAPJ.demos@WMQSOAP.DEMO.QM&"
 + "connectionFactory=connectQueueManager(WMQSOAP.DEMO.QM)"
 + "&initialContextFactory=com.ibm.mq.jms.NoJndi&targetService=soap.server.StockQuoteAxis.java";

 QName serviceName = new QName("soap.server.StockQuoteAxis_Wmq", "StockQuoteAxisService");
 QName portName = new QName("soap.server.StockQuoteAxis_Wmq", "soap.server.StockQuoteAxis_Wmq");

 Service service = Service.create(serviceName);
 service.addPort(portName, SOAPBinding.SOAP11HTTP_BINDING, endpointUrl);

 /** Create a Dispatch instance from a service.*/
 System.out.println("Creating dispatch instance for service StockQuoteAxisService");
 Dispatch<SOAPMessage> dispatch = service.createDispatch(portName, SOAPMessage.class,
 Service.Mode.MESSAGE);
 System.out.println(" > Dispatch instance created.");

 /** Create SOAPMessage request. */
 MessageFactory mf = MessageFactory.newInstance(SOAPConstants.SOAP_1_1_PROTOCOL);

 System.out.println("Creating a Request Reply SOAP Message");
 SOAPMessage request = mf.createMessage();

 /** Obtain the SOAPEnvelope and header and body elements */
 SOAPPart part = request.getSOAPPart();
 SOAPEnvelope env = part.getEnvelope();
 SOAPHeader header = env.getHeader();
 SOAPBody body = env.getBody();

 /** Construct the message payload. */
 SOAPElement operation = body.addChildElement("getQuote", "ns1",
 "soap.server.StockQuoteAxis_Wmq");
 SOAPElement value = operation.addChildElement("in0");
 value.addAttribute(new QName("http://www.w3.org/2001/XMLSchema-instance", "type"), "string");
 value.addTextNode("XXX");
 request.saveChanges();
 System.out.println(" > SOAP Message created.");

 /** Invoke the service endpoint. */
 DispatchClientAsyncCallback handler = new DispatchClientAsyncCallback();

 System.out
 .println("Invoking getQuote Request Reply operation asynchronously using a callback...");
 Future<> monitor = dispatch.invokeAsync(request, handler);
 System.out.println(" > getQuote call has returned");

 /** Sleep main thread until handler has been notified */
 System.out.println("Waiting for handler to be called...");

```

## Developing a .NET 1 or 2 client for WebSphere transport for SOAP using Microsoft Visual Studio 2012:

Develop an .NET 1 or 2 web service client to run using the IBM MQ transport for SOAP.

### Before you begin

**Important:** .NET Framework 2.0 is a prerequisite for this task, therefore make sure that you have it installed before you begin.

You can start the development of a .NET 1 or 2 client in a number of different ways:

1. Use **amqwdeployWmqService** to generate client stubs from a web service and import them into Microsoft Visual Studio.
2. Use **java2wsdl** to generate WSDL from a Java implementation of a web service, and then use **wsdl.exe**, which is shipped with .NET, to generate client stubs.
3. Generate WSDL from a .NET **.asmx** implementation of the service using **amqswsdl**, and then use **wsdl.exe**.
4. If you have developed and deployed the service for HTTP, use the **Add Web reference...** wizard in Microsoft Visual Studio to configure the client to access the HTTP service. Alter the URL refer to the service deployed to IBM MQ transport for SOAP.

The task uses the service developed in “Developing a .NET 1 or 2 service for IBM MQ transport for SOAP using Microsoft Visual Studio 2012” on page 1233.

### About this task

Follow these steps to create a .NET 1 or 2 Client for HTTP and IBM MQ transport for SOAP.

### Procedure

1. Create the client console application and modify it to invoke the StockQuote HTTP web service.
  - a. Right click **Solution 'StockQuoteDotNet'** in the Solution Explorer > Add... > New Project. Select the **C# Project Type, .NET Framework 2.0, and Console Application**. Name the project **StockQuoteClientDotNet** > **OK**
  - b. Right click **Solution 'StockQuoteDotNet'** in the Solution Explorer > Add... > New Project. Select the **C# Project Type, .NET Framework 2.0, and Console Application**. Name the project **StockQuoteClientDotNet** > **OK**
  - c. Right-click **StockQuoteClientDotNet** > **Set as Startup project**.
  - d. Right-click **StockQuoteClientDotNet** > **Add Web Reference...** > Browse to web services in this solution > Select **StockQuote** > **Add Reference**. Notice you have added a Web reference to local host and a new configuration file **app.config**.
  - e. In the Solution Explorer, change the name of the console application from **Program.cs** to **StockQuoteClientDotNet.cs** > Click **OK** to changing all the usages of **Program.cs** to **StockQuoteClientDotNet.cs**.
  - f. Replace the contents of **StockQuoteClientDotNet.cs** with the code in Figure 178 on page 1267.

---

```

using System;
using StockQuoteClientDotNet.localhost;
namespace StockQuoteClientDotNet {
 class StockQuoteClientDotNet {
 static void Main(string[] args) {
 try {
 StockQuote stockobj = new StockQuote();
 Console.WriteLine("http reply is: "
 + stockobj.getNonInlineQuote("http request"));
 }
 catch (System.Exception e) {
 Console.WriteLine("Exception thrown: " + e.ToString());
 }
 Console.ReadLine();
 }
 }
}

```

---

*Figure 178. HTTP StockQuoteClientDotNet program*

g. Launch StockQuoteClientDotNet to test against the StockQuote.asmx service:

1) Press **F5**, click the green arrow in the action bar, or **Debug > Start Debugging (F5)**.

If the StockQuoteDotNet project is in the same solution, it starts automatically. Otherwise you need to start the service first.

The command window with the results opens behind the workspace. The Console.ReadLine(); statement prevents it from closing until you press **Enter**.

**Tip:** Make sure StockQuote.asmx is the Start page in the StockQuoteDotNet project.

2. Modify StockQuoteClientDotNet to call the StockQuote.asmx service using IBM MQ transport for SOAP.

a. Add the lines shown in bold to the client.

---

```

using System;
using StockQuoteClientDotNet.localhost;
namespace StockQuoteClientDotNet {
 class StockQuoteClientDotNet {
 static void Main(string[] args) {
 try {
 IBM.WMQSOAP.Register.Extension();
 StockQuote stockobj = new StockQuote();
 Console.WriteLine("http reply is: "
 + stockobj.getNonInlineQuote("http request"));
 stockobj.Url = "jms:/queue?"
 + "initialContextFactory=com.ibm.mq.jms.Nojndi"
 + "&connectionFactory=()&destination=REQUESTDOTNET@QM1"
 + "&targetService=StockQuote.asmx";
 Console.WriteLine("jms reply is: "
 + stockobj.getNonInlineQuote("jms request"));
 } catch (System.Exception e) {
 Console.WriteLine("Exception thrown: " + e.ToString());
 }
 Console.ReadLine();
 }
 }
}

```

---

Figure 179. Modified StockQuoteClientDotNet program

Alternatively, modify the default URL. Open **StockQuoteClientDotNet > Properties > Settings.settings** and change the value of the `StockQuoteClientDotNet_localhost_StockQuote` property to the IBM MQ transport for SOAP URL.

- b. Add a reference to `amqsoap.dll`
  - 1) In the **StockQuoteClientDotNet** project in the Solution Explorer, right-click **References > Add Reference...** > Click the **Browse** tab > browse to `MQ_INSTALLATION_PATH\bin` > Select **amqsoap.dll** > **OK**. `MQ_INSTALLATION_PATH` is the directory where IBM MQ is installed.
3. Test the client with the `StockQuote.asmx` service using IBM MQ transport for SOAP.
  - a. Open a command window in the `StockQuoteDotNet` project directory: `.\StockQuoteDotNet\StockQuoteDotNet > Verify that .bin\StockQuoteDotNet.dll exists. If it does not, rebuild the solution.`
  - b. Type the command **amqwRegisterdotNet**. You need only run **amqwRegisterdotNet** once per installation.
  - c. If you have run **amqwdeployWMQServer** with the `genAsmxWMQBits`, run the `.NET SOAP Listener`: `generated\server\startWMQNListener`
  - d. Alternatively run the listener directly:
 

```

amqwSOAPNETListener -u "jms:/queue?
destination=REQUESTDOTNET@QM1
&connectionFactory=()&initialContextFactory=com.ibm.mq.jms.Nojndi
&targetService=StockQuote.asmx&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE"
-w C:\IBM\ID\StockQuoteDotNet\StockQuoteDotNet -n 10

```
4. In Visual Studio 2012, press **F5** to run `StockQuoteClientDotNet`.

### .NET Framework 1 and .NET Framework 2 web service clients

The sample .NET clients provided with the IBM MQ transport for SOAP use generated stubs to call the sample Axis and .NET services.

For .NET Framework 1 and .NET Framework 2 clients, IBM MQ provides access to web services using .NET clients. The **amqwdeployWMQService** command has an option, `genProxiestoDotNet`, that generates

.NET Framework 1 or .NET Framework 2 client stubs for a web service. You can also use client stubs generated by the .NET **wsdl** tool, or by Microsoft Visual Studio 2012.

The sample .NET Framework 1 and .NET web service clients are installed in `MQ_INSTALLATION_PATH\tools\soap\samples\dotnet`. `MQ_INSTALLATION_PATH` is the directory where IBM MQ is installed.

#### **SQVB2Axis.vb**

`SQVB2Axis.vb`, Figure 180, is the Visual Basic client to call the **StockQuoteAxisService** service.

#### **SQVB2DotNet.vb**

`SQVB2DotNet.vb`, Figure 181, is the Visual Basic client to call the **StockQuoteDotNet** service.

#### **SQCS2Axis.cs**

`SQCS2Axis.cs`, Figure 182 on page 1270, is the C# client to call the **StockQuoteAxisService** service. You can override the URL of the service by providing a URL on the command line.

#### **SQCS2DotNet.cs**

`SQCS2DotNet.cs`, Figure 183 on page 1270, is the C# client to call the **StockQuoteDotNet** service. You can override the URL of the service by providing a URL on the command line.

---

```
Module SQVB2Axis
 Function Main(ByVal CmdArgs() As String) As Integer
 IBM.WMQSOAP.Register.Extension()
 Dim obj As New StockQuoteAxisService()
 Dim res As Single = obj.getQuote("fromcs")
 System.Console.WriteLine("SQVB2Axis: reply is: '{0}'", res)
 End Function
End Module
```

---

*Figure 180. SQVB2Axis*

---

```
Module SQVB2DotNet
 Function Main(ByVal CmdArgs() As String) As Integer
 IBM.WMQSOAP.Register.Extension()
 Dim obj as new StockQuoteDotNet()
 Dim res as Single = obj.getQuote("fromcs")
 System.Console.WriteLine("SQVB2DotNet: reply is: '{0}'", res)
 End Function
End Module
```

---

*Figure 181. SQVB2DotNet*

---

---

```

using System;
class SQCS2Axis {
 [STAThread]
 static void Main(string[] args) {
 try {
 IBM.WMQSOAP.Register.Extension();
 StockQuoteAxisService stockobj = new StockQuoteAxisService();
 if (args.GetLength(0) >= 1)
 stockobj.Url = args[0];
 System.Single res = stockobj.getQuote("XXX");
 Console.WriteLine("SQCS2Axis RPC reply is: " + res);
 }
 catch (System.Exception e) {
 Console.WriteLine("\n>>> EXCEPTION WHILE RUNNING SQCS2Axis DEMO <<<\n"
 + e.ToString());
 }
 }
}

```

---

*Figure 182. SQCS2Axis*

---

```

using System;
class SQCS2DotNet {
 [STAThread]
 static void Main(string[] args) {
 try {
 IBM.WMQSOAP.Register.Extension();
 StockQuoteDotNet stockobj = new StockQuoteDotNet();
 if (args.GetLength(0) >= 1)
 stockobj.Url = args[0];
 System.Single res = stockobj.getQuote("XXX");
 Console.WriteLine("RPC reply is: " + res);
 if (args.GetLength(0) == 0) {
 res = stockobj.getQuoteDOC("XXX");
 Console.WriteLine("DOC reply is: " + res);
 }
 }
 catch (System.Exception e) {
 Console.WriteLine("\n>>> EXCEPTION WHILE RUNNING SQCS2DotNet DEMO <<<\n"
 + e.ToString());
 }
 }
}

```

---

*Figure 183. SQCS2DotNet*

## Deploying web services using the IBM MQ transport for SOAP

Deploy a web service to one of a number of different server environments and connect to it using IBM MQ transport for SOAP.

### Before you begin

Develop a web service and test it using SOAP over HTTP in the target environment.

### About this task

You can deploy a web service to run with IBM MQ transport for SOAP in a number of different SOAP run time environments. You can deploy a service to Axis 1.4 using only the software installed with IBM MQ. For the other run time environments, you must install additional software.

You are not restricted to running IBM MQ transport for SOAP to the servers for which there are deployment instructions. Use the instructions to deploy a service to one of the listed environments.

**Note:** Some integrated environments offer SOAP over JMS using the W3C recommended JMS SOAP binding, as well as the IBM MQ transport for SOAP binding. Releases of IBM MQ, up to and including 7.0.1.2, support only the IBM MQ transport for SOAP binding. From 7.0.1.3 onwards you can deploy Axis2 clients using a URI that conforms to the W3C candidate recommendation for SOAP over JMS. See the tutorial, *Develop a SOAP/JMS JAX-WS web services application with WebSphere Application Server V7 and Rational Application Developer V7.5*.

### Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using `amqwdeployWMQService`:

Deploy an Axis 1.4 service to IBM MQ transport for SOAP by creating a deployment directory, running the `amqwdeployWMQService` command, and starting the Axis 1.4 listener.

### Before you begin

1. Follow the instructions for installing IBM MQ transport for SOAP
2. Verify the installation and your environment using the `runivt` command.
3. To redeploy a service:
  - a. Delete the `./generated` subdirectory, and all its subdirectories.
  - b. Remove requests from the destination queue and delete it.
  - c. Proceed with the instructions from step 2.

### About this task

These instructions are to deploy an Axis 1.4 service for the first time. To restart an Axis 1.4 service, rerun the Axis 1.4 SOAP listener: step 11 on page 1272.

Use the following instructions to deploy a new Axis 1.4 service to IBM MQ transport for SOAP:

### Procedure

1. Create a directory `deployDir` to hold the deployment files. The deployment utility requires that each service is deployed from a separate directory.
2. Open a command window on Windows, or a command shell using X Window System on UNIX and Linux systems, in `deployDir` to run `amqwdeployWMQService`.
3. Run `amqwsetcp` to set the classpath. The JRE and JDK must be in the classpath, at version 5.0 or later, and at the same version level.
4. Copy the class source, `className.java`, into `deployDir`

5. Copy all the Java source files in the same package as *className* into *deployDir/packageName*, where *packageName* is a directory tree corresponding to the package name.
6. Run **javac** *packageName.className*. You might need to add a path to the current directory ". ", or to the *packageName* directory for **javac** to find the other classes.

7. Create the Axis WSDL for the service:

```
amqwdeployWMQService -f packageName.className.java -c genAxisWsd1
-v -u "jms:/queue?destination=queueName
&initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=(connectQueueManager(QmgrName)binding(auto))"
```

8. Create the IBM MQ resources for the service:

```
amqwdeployWMQService -f packageName.className.java -c genAxisWMQBits
-v -u "jms:/queue?destination=queueName
&initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=(connectQueueManager(QmgrName)binding(auto))"
```

**Tip:**

If you want to set up a new queue manager, and the resources it needs, to do development and testing, run **setupWMQSOAP**.

If you want set up the new queue manager as the default, take a copy of **setupWMQSOAP** from the *WMQ installation directory\tools\soap\samples* directory, and add the -q parameter to the line call :try -q crtmqm %QMGR%

9. Create the Axis listener and deploy the service:

```
amqwdeployWMQService -f packageName.className.java -c AxisDeploy
-v -u "jms:/queue?destination=queueName
&initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=(connectQueueManager(QmgrName)binding(auto))"
```

10. If you need to generate the WSDL for the service, generate client stubs, or client proxies, run **amqwdeployWMQService** with one of the following parameters:

- genAsmxWsd1
- genAxisWsd1
- genProxiesToDotNet
- genProxiestoAxis

**Note:** You must generate WSDL before generating the proxies. The AllAxis option fails if the CLASSPATH is not set up to find all the classes that are imported to compile *className.java*. If there are multiple Java files in the package containing *className.java*, you must compile them first using **javac**. **amqwdeployWMQService -f *packageName.className.java* -c CompileJava** compiles only *className.java*.

11. Start the generated Axis listener.
 

```
.\generated\server\startWMQJListener.cmd
```



## Deploying a service to .NET Framework 1 or 2 service to use IBM MQ transport for SOAP:

Deploy a .NET Framework 1 or 2 service to IBM MQ transport for SOAP. Create a deployment directory, run the `amqwdeployWMQService` command, and start the .NET listener.

### Before you begin

1. Follow the instructions for installing IBM MQ transport for SOAP
2. Verify the installation and your environment using the `runivt` command.
3. The path to the .NET framework files `wsdl.exe` and `csc.exe` must be set. The copies of `wsdl.exe` and `csc.exe` identified by the `PATH` variable must be at the same level of the .NET framework. If you have multiple .NET frameworks installed, or are using Visual Studio, check the `PATH` variable carefully.
4. To redeploy a service:
  - a. Delete the `./generated` subdirectory, and all its subdirectories
  - b. Remove requests from the destination queue and delete it.
  - c. Proceed with the instructions from step 2.

### About this task

These instructions are to deploy a .NET service for the first time. To restart a .NET service, rerun the .NET SOAP listener, step 9 on page 1274.

Use the following instructions to deploy a new .NET Framework 1 or .NET Framework 2 service to IBM MQ transport for SOAP:

### Procedure

1. Create a directory `deployDir` to hold the deployment files. The deployment utility requires that each service is deployed from a separate directory.
2. Open a command window in `deployDir` to run `amqwdeployWMQService`.

```
C:\IBM\ID\QuoteClient>
```

3. Run `amqwsetcp` to set the classpath. A classpath is needed only for Axis clients.
4. Copy the .NET service, `className.asmx`, into `deployDir`
5. Build the service implementation into a library ( `.dll` ).

The inline service implementation is in `className.asmx`. The code-behind service implementation might be `className.asmx.cs`.

Figure 184 shows an example of a command to build a .NET Framework V2 service as a library.

---

```
c:\WINDOWS\Microsoft.NET\Framework\v3.5\Csc.exe /noconfig /nowarn:1701,1702
/errorreport:prompt /warn:4 /define:TRACE
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.configuration.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Data.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Drawing.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Web.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Web.Services.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Xml.dll
/debug:pdbonly /filealign:512 /optimize+
/out:obj\Quote.dll /target:library Properties\AssemblyInfo.cs Quote.asmx.cs
```

---

Figure 184. Build command for .NET Framework V2 service

6. Copy `className.dll` into `deployDir\bin`.

7. Set up the IBM MQ resources, and create the listener required for the service:

```
amqwdeployWMQService -f className.asmx -c genAsmxWMQBits
-v -u "jms:/queue?destination=queueName
&initialContextFactory=com.ibm.mq.jms.NoJndi
&connectionFactory=(connectQueueManager(QmgrName)binding(auto))
&targetService=className.asmx"
```

8. If you need to generate the WSDL for the service, generate client stubs, or client proxies, run **amqwdeployWMQService** with one of the following parameters:

- genAsmxWsd1
- genAxisWsd1
- genProxiesToDotNet
- genProxiestoAxis

**Note:** You must generate WSDL before generating the proxies.

9. Start the generated .NET listener.

```
.\generated\server\startWMQNListener.cmd
```

### Deploying a service to CICS Transaction Server to use WebSphere Transport for SOAP:

IBM MQ transport for SOAP is integrated into CICS Transaction Server 4.1 web services support.

#### Before you begin

Use the same tools to develop for a client or service for IBM MQ, as you would to develop for HTTP. CICS has tools corresponding to **Java2wsdl** and **wsdl2Java**:

- **DFHWS2LS** takes a web service description as a starting point. It uses the descriptions of the messages, and the data types used in those messages, to construct high-level language data structures. You can use in the structures in application programs written in different languages.
- **DFHLS2WS** takes a high-level language data structure as a starting point. It uses the structure to construct a web services description that contains descriptions of messages. It also creates schemas for the messages from the language data structure.

Follow the Creating a web service instructions in the CICS product documentation, to create a web service.

#### About this task

Follow the instructions, Configuring CICS to use the IBM MQ transport in the CICS product documentation. Using the instructions, you can deploy the web service to IBM MQ transport for SOAP.

## **Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP:**

IBM MQ transport for SOAP is integrated into the service integration bus on WebSphere Application Server.

### **Before you begin**

Use Rational Application Developer, WebSphere Integration Developer, or a web services toolkit to develop the web service.

### **About this task**

Use the following instructions to deploy a service using IBM MQ transport for SOAP as a SOAP transport on WebSphere Application Server.

### **Procedure**

1. Configure IBM MQ as the JMS messaging provider for the service integration bus on WebSphere Application Server.
2. Configure the IBM MQ resources required by the service.
3. Follow the instructions, *Configuring JMS resources for the synchronous SOAP over JMS endpoint listener*, in the WebSphere Application Server Network Deployment product documentation. There are corresponding instructions for other WebSphere Application Server platforms.
4. Modify the service URI to conform to the IBM MQ transport for SOAP URI.
5. Deploy the service to WebSphere Application Server.

### **What to do next**

Deploy the service with HTTP as a transport so that clients can query the service and receive the WSDL in response.

## **Deploying a service to WebSphere ESB and Process Server service endpoint to use WebSphere Transport for SOAP:**

IBM MQ transport for SOAP is not directly supported by WebSphere ESB and Process Server. You must configure a custom Export.

### **About this task**

WebSphere Integration Developer provides a SOAP data transformation that you can bind to the IBM MQ JMS Export to create a custom IBM MQ JMS SOAP Export.

Follow the instructions to create a customized Export to receive SOAP requests over IBM MQ transport for SOAP.

### **Procedure**

1. Read “Overview of imports and exports” and “How to connect to IBM MQ” in the WebSphere Process Server for Multiplatforms product documentation. See *Overview of imports and exports* and *How to connect to IBM MQ*.
2. Follow the task “Generating an MQ JMS export binding” in the WebSphere Integration Developer product documentation. See *Generating an MQ JMS export binding*. Use the prepackaged SOAP data-binding described in *Prepackaged JMS data format transformations* to format the SOAP message.

## Deploying web service clients to use IBM MQ transport for SOAP

Deploy a web service client to one of a number of different client environments and connect to a service using IBM MQ transport for SOAP.

### Before you begin

Develop the web service and deploy it to use IBM MQ transport for SOAP.

### About this task

You can deploy a web service client to run with IBM MQ transport for SOAP in a number of different client environments. You can deploy a Java client to Axis 1.4 using only the software installed with IBM MQ. For the other client environments, you must install additional software.

You are not restricted to running WebSphere transport for SOAP in the client environments for which there are deployment instructions. Use the instructions to deploy a client to one of the supported environments.

**Note:** Some integrated environments offer SOAP over JMS using the W3C recommended JMS SOAP binding, as well as the IBM MQ transport for SOAP binding. Releases of IBM MQ, up to and including 7.0.1.2, support only the IBM MQ transport for SOAP binding. From 7.0.1.3 onwards you can deploy Axis2 clients using a URI that conforms to the W3C candidate recommendation for SOAP over JMS. See the tutorial, *Develop a SOAP/JMS JAX-WS web services application with WebSphere Application Server V7 and Rational Application Developer V7.5*.

### Deploying a web service client to Axis 1.4 to use IBM MQ transport for SOAP:

Prepare a deployment directory and deployment descriptor for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure IBM MQ queues and channels, start the service and test the client.

### Before you begin

**Tip:** Deploy the service to HTTP, develop and test the client for HTTP, and then modify the client for IBM MQ transport for SOAP:

1. Add the `Register.extension()` call to the client.
2. Change the static web service address in the client proxy locator class to use the URI for the IBM MQ transport for SOAP.

### About this task

Deploying an Axis 1.4 client to use IBM MQ transport for SOAP requires one additional deployment step compared to an HTTP client. You must create a client deployment descriptor, `client-config.wsdd`, to map the `jms:` transport to the sender class `com.ibm.mq.soap.transport.jms.WMQSender`.

If you use the command `amqwdeployWMQService` to generate client proxies, you can deploy the client using the directories the command generates.

### Procedure

1. Create a directory `deployDir` to hold the client deployment files.
2. Open a command window on Windows systems, or a command shell using X Window System on UNIX and Linux systems, in `deployDir`.
3. Run the `amqwsetcp.cmd` command to set the CLASSPATH

4. Run the **amqwclientconfig.cmd** command to create an Axis 1.4 client deployment descriptor, `client-config.wsdd` in *deployDir*.
5. Make sure the classes in the client package, the client proxy classes, and the libraries the client uses, are in the CLASSPATH.  
**amqwdeployWMQService** places the .NET client proxies into `./generated/server/soap/client/remote/dotnetService` and the Axis 1.4 proxies into `./generated/server/soap/client/remote/client package`.

## Example

The example shows the configuration and output, Figure 187 on page 1278, from an Axis 1.4 Java client. The client, Figure 186, calls a web service that echoes its input parameter. The service definition, Figure 185, shows the URI taken from the service WSDL.

---

```
<wsdl:service name="QuoteSOAPImplService">
 wsdl:port binding="intf:org.example.www.QuoteSOAPImplBindingSoap"
 name="org.example.www.QuoteSOAPImpl_Wmq">
 <wsdlsoap:address location="jms:/queue?destination=REQUESTAXIS
 &connectionFactory=(connectQueueManager(QM1)binding(server))
 &initialContextFactory=com.ibm.mq.jms.NoJndi
 &targetService=org.example.www.QuoteSOAPImpl.java
 &replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE" />
 </wsdl:port>
</wsdl:service>
```

---

Figure 185. Service definition

---

```
package org.example.www;
import com.ibm.mq.soap.Register;
public class QuoteClient {
 public static void main(String[] args) {
 try {
 Register.extension();
 QuoteSOAPImplServiceLocator locator = new QuoteSOAPImplServiceLocator();
 System.out.println("Response = "
 + locator.getOrgExampleWwwQuoteSOAPImpl_Wmq().getQuote("IBM"));
 } catch (Exception e) {
 System.out.println("Exception = " + e.getMessage());
 }
 }
}
```

---

Figure 186. Axis 1.4 Java client

```

C:\IBM\ID\Test>dir /s /b
C:\IBM\ID\Test\client-config.wsdd
C:\IBM\ID\Test\org
C:\IBM\ID\Test\org\example
C:\IBM\ID\Test\org\example\www
C:\IBM\ID\Test\org\example\www\GetQuoteFaultMsg.class
C:\IBM\ID\Test\org\example\www\OrgExampleWwwQuoteSOAPImplBindingSoapStub.class
C:\IBM\ID\Test\org\example\www\QuoteClient.class
C:\IBM\ID\Test\org\example\www\QuoteSOAPImpl.class
C:\IBM\ID\Test\org\example\www\QuoteSOAPImplService.class
C:\IBM\ID\Test\org\example\www\QuoteSOAPImplServiceLocator.class

C:\IBM\ID\Test>amqwsetcp
C:\IBM\ID\Test>java org.example.www.QuoteClient.class
Response = IBM

```

Figure 187. Client configuration and output

### What to do next

1. If you are deploying the client as an IBM MQ client, configure the client and server connection channel.
2. If you are deploying the client to a different queue manager to the service, you must make the destination queue available to the client. Configure the destination queue on the service queue manager as a cluster queue, or on the client queue manager as a remote queue definition.

### Deploying a web service client to Axis2 to use IBM MQ transport for SOAP:

Prepare a deployment directory and Axis2 configuration file for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure IBM MQ queues and channels, start the service and test the client.

### Before you begin

**Tip:** Deploy the service to HTTP. Develop and test the client for HTTP, and then modify the URL to reference the service using IBM MQ transport for SOAP.

The task shows how to deploy an unmanaged Axis2 client to Java Standard Edition. You might want to deploy an Axis2 client to a Web container. In “Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 1249, you developed a client in a Web container and deployed it to WebSphere Application Server Community Edition. As part of the server configuration, you enabled the Axis2 facet and included the facet in the configuration of the Web container. To configure Web containers on other application servers, refer to the Axis2 documentation, [http://ws.apache.org/axis2/1\\_4\\_1/installationguide.html#servlet\\_container](http://ws.apache.org/axis2/1_4_1/installationguide.html#servlet_container), or the documentation supplied with the Web server.

**Note:** Axis2 use the term, Servlet container. A Servlet container is the same as a Web container.

### About this task

Deploying an Axis2 client to use IBM MQ transport for SOAP is like deploying an Axis2 client to use HTTP. Additional steps are required to provide a classpath to the IBM MQ JAR files, and to modify the Axis2 configuration file. The Axis2 configuration file requires an additional entry for JMS. The entry refers to the IBM MQ transport for SOAP JAR file that implements the JMS transportSender.

Axis2 provides a script, `axis2.bat` or `axis2.sh`, which simplifies client deployment; see the examples in Figure 191 on page 1281 and Figure 192 on page 1281.

### Note:

1. `axis2.bat` has a bug that must be corrected. The string `-Djava.ext.dirs="%AXIS2_HOME%\lib\"` must be changed to `-Djava.ext.dirs="%AXIS2_HOME%\lib\\"`.

2. In `axis2.bat` and `axis2.sh`, `-Djava.ext.dirs` is used as a quick way to reference all the Axis2 JAR files, instead of adding them separately to the classpath. Unfortunately this approach is flawed, and only works with some JREs. It does not work with the IBM JREs.

The JVM parameter, `-Djava.ext.dirs="%AXIS2_HOME%\lib\\"`, makes the Axis JAR files available to the JVM. The JVM attempts to instantiate some of the Axis JAR files, and leads to an error, the details of which depend on the JVM. Typically, you might see one of the following lines in the stack trace:

```
org.apache.axiom.om.util.UUIDGenerator.getInitialUUID(UUIDGenerator.java:76)
```

```
or org.apache.axis2.deployment.DeploymentException: java.security.NoSuchAlgorithmException:
MD5 MessageDigest not available
```

The correct way to run an unmanaged Axis2 client is to add the Axis2 JAR files to the classpath. The classpath is available only to the client application and not to the JVM.

The procedure describes the general steps to run an unmanaged Axis2 client without using the `axis2` script. The examples in Figure 189 on page 1280 and Figure 190 on page 1280 are scripts for Windows and Linux.

### Procedure

1. Download Axis2 1.4.1 from [http://ws.apache.org/axis2/download/1\\_4\\_1/download.cgi](http://ws.apache.org/axis2/download/1_4_1/download.cgi) and unpack into a folder, `Axis2-1.4.1`.
2. Update `axis2.xml` in `Axis2-1.4.1\conf`.

- a. Update `axis2.xml` in `Axis2-1.4.1\conf`. Add IBM MQ transport for SOAP as a `transportSender`:

```
<transportSender name="jms"
class="com.ibm.mq.axis2.transport.jms.WMQJMSTransportSender"/>
```

- b. If required, alter the size of the connection pool from the default of 10.

```
<transportSender name="jms"
class="com.ibm.mq.axis2.transport.jms.WMQJMSTransportSender">
<parameter name="ResourcePoolCapacity">20</parameter>
</transportSender>
```

`ResourcePoolCapacity` defines how many service endpoint entries are kept in the cache. The value must be at least 1. If the number of service endpoint entries exceeds the cache size, entries are deleted to make room for new entries. The size of an endpoint entry varies. Set a number that is large enough to avoid the cache thrashing.

See step 3 in “Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 1249.

3. Create a directory `deployDir`. Under this directory copy the folder structure containing the client and client proxies. `deployDir` is equivalent to the `project\bin` folder in an Eclipse Java project.
4. Open a command window on Windows, or a command shell using X Window System on UNIX and Linux systems, in `deployDir`.
5. Update the classpath to include the current directory, Axis2 JAR files, `com.ibm.mqjms.jar` and `com.ibm.mq.axis2.jar`. `com.ibm.mqjms.jar` references all the other IBM MQ JAR files that are required.
6. Use the **Java** command to start the client program.

### Examples

Four example of running an Axis2 client are listed in Figure 190 on page 1280 to Figure 192 on page 1281. Figure 188 on page 1280 shows the output from running the asynchronous client listed in Figure 170 on page 1255.

```

cd C:\IBM\ID\Workspaces\Axis2docs\StockQuoteAxis2PojoClient\bin>
runpojo soap/client/SQA2AsyncClient

HTTP Sync: 55.25
Callback constructor
Waiting for HTTP callback
Result in Callback 55.25
HTTP poll: 55.25
JMS: Sync: 55.25
Callback constructor
Waiting for JMS callback
Result in Callback 55.25
JMS poll: 55.25
Press any key to continue . . .

```

Figure 188. Output from running SQA2AsyncClient

```

@echo off
set AXIS2_HOME=C:\OpenSource\axis2-1.4.1
set JAVA_HOME=C:\IBM\Java50
set WMQ_HOME=C:\IBM\MQ\java\lib

setlocal EnableDelayedExpansion
set CLASSPATH=
set AXIS2_CLASS_PATH=
FOR %%c in ("%AXIS2_HOME%\lib*.jar") DO set AXIS2_CLASS_PATH=!AXIS2_CLASS_PATH!;%%c

"%JAVA_HOME%\bin\java" -Daxis2.repo="%AXIS2_HOME%\repository"
-Daxis2.xml="%AXIS2_HOME%\conf\axis2.xml" -cp
";%WMQ_HOME%\com.ibm.mqjms.jar;%WMQ_HOME%\com.ibm.mq.axis2.jar;%AXIS2_CLASS_PATH%" %1

pause

```

Figure 189. runpojo.bat: Windows, using a classpath

```

export AXIS2_HOME=/home/OpenSource/axis2-1.4.1
export JAVA_HOME=/usr/lib/j2sdk1.5-ibm
update classpath
AXIS2_CLASSPATH=""
for f in "$AXIS2_HOME"/lib/*.jar
do
 AXIS2_CLASSPATH="$AXIS2_CLASSPATH":$f
done
AXIS2_CLASSPATH="$AXIS2_HOME": "$JAVA_HOME/lib/tools.jar": "$AXIS2_CLASSPATH": "$CLASSPATH"
java -cp /home/alex/dev/sandbox/Soap/axis2:/opt/mqm/java/lib/com.ibm.mqjms.jar:
/opt/mqm/java/lib/com.ibm.mq.axis2.jar:$AXIS2_CLASSPATH
-Daxis2.xml=/home/alex/dev/sandbox/axis2-1.4.1/conf/axis2.xml %1

```

Figure 190. runpojo.sh: Linux, using a classpath.



---

```
@echo off
set AXIS2_HOME=C:\OpenSource\axis2-1.4.1
set JAVA_HOME=C:\IBM\Java50
set WMQ_HOME=C:\IBM\MQ\java\lib

%AXIS2_HOME%\bin\axis2 -cp .;%WMQ_HOME%\com.ibm.mqjms.jar;%WMQ_HOME%\com.ibm.mq.axis2.jar; %1
pause
```

---

*Figure 191. runaxis2.bat: Windows, using axis2.bat*

#### Note

---

```
export AXIS2_HOME=/home/OpenSource/axis2-1.4.1
export JAVA_HOME=/usr/lib/j2sdk1.5-ibm

%AXIS2_HOME%\bin\axis2 -cp .;%WMQ_HOME%\com.ibm.mqjms.jar;%WMQ_HOME%\com.ibm.mq.axis2.jar; %1
```

---

*Figure 192. runaxis2.sh: Linux, using axis2.sh*

#### Note

### Deploying to an Axis2 client using W3C SOAP over JMS:

A web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 4 of connecting an Axis2 web service client and a web service deployed to WebSphere Application Server using the W3C SOAP over JMS protocol. Modify the URL in the Axis2 client developed for IBM MQ transport for SOAP to use the W3C candidate recommendation for SOAP over JMS.

#### Before you begin

You must first complete the task, “Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 1249 to call **SimpleJavaListener** using an Axis2 client and the IBM MQ transport for SOAP protocol.

You must also have created the web service and configured IBM MQ and WebSphere Application Server.

See Planning to use IBM MQ and WebSphere Application Server together.

You must also have completed the following task, “Developing a JAX-WS EJB web service for W3C SOAP over JMS” on page 1236.

In the task, the client runs in Eclipse Galileo. You might run the client from the command line by modifying the `Axis2.bat` file shipped with Axis2.

#### About this task

The only change you must make to the existing Axis2 `StockQuoteAxis` static client to call the `StockQuoteAxis` service hosted by WebSphere Application Server is to change the URL passed to the client. Since the WSDL has not changed, you can use the same proxy classes in the `soap.server` package.

You have two approaches to defining the URL to pass to the client. You might use the same URL as in the generated `StockQuoteAxis.wsdl`. You must add the `jndiInitialContextFactory` and the `jndiURL`

parameters to access the WebSphere Application Server JNDI directory. Another approach is to change the URL and give the client direct access to the REQUESTAXIS and REPLYAXIS queues on QM1, without using a JNDI lookup.

The connection parameters you define in the URL passed to the Axis2 client are used to connect to the IBM MQ queue manager and queues required to send and receive SOAP messages. The connection parameters passed to the Axis2 client are not necessarily used by the service. You can use the distributed queuing capabilities of IBM MQ to decouple the client and service from using the same queue manager, or the same name server.

### Procedure

1. Save the URL from the generated `StockQuoteAxis.wsdl` and close down Rational Application Developer to save on memory.

If you did not change the server configuration, closing Rational Application Developer stops the application server. In which case start the server with the command:

```
startserver server1
```

2. Open Eclipse Galileo in the workspace with the Axis2 client project.
3. Open `SQA2StaticClient.java`.  
See `SQA2StaticClient.java`.
4. Call the service using the queue variant of the URI.

- a. Modify the URL. The new URI is:

```
jms:queue:REQUESTAXIS
 ?replyToName=REPLYAXIS
 &connectionFactory=connectQueueManager(QM1)Bind(Server)
 &targetService=StockQuoteAxis;
```

Compare this to the URL from `StockQuoteAxis.wsdl`:

```
jms:jndi:requestaxis
 ?jndiConnectionFactoryName=qm1
 &targetService=StockQuoteAxis
```

*Figure 193. URL from StockQuoteAxis.wsdl*

- REQUESTAXIS is now uppercase as it is a queue name and not a JNDI name.
  - The connection to QM1 is straightforward.
  - The URI does not contain the name of the Reply to destination. The client must define the queue it expects replies on.
- b. Run `SQA2StaticClient.java` using the same Run as... configuration as you did in the task, "Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse" on page 1249.
5. Call the service using the jndi variant of the URI, using WebSphere Application Server as the naming server.
    - a. Use the URL from `StockQuoteAxis.wsdl`, Figure 193, providing the missing parameters to use the naming service in WebSphere Application Server.  
The missing parameters and values you must provide are:

Table 185. Additional JNDI parameters

Parameter	Value used in this example	Description
&jndiURL	iop://localhost:2810 or corbaname:iop:localhost:2810	URI of naming provider. For WebSphere Application Server the value defaults to 2809. It is also known as the port number of the RMI connector, and the bootstrap port. The value is listed in the SystemOut.log  00000000 NameServerImp A NMSV0018I: Name server available on bootstrap port 2810
&jndiInitialContextFactory	com.ibm.websphere.naming.WsnInitialContextFactory	The name of the initial context factory used by WebSphere Application Server.
&replyToName	replyaxis	JNDI name of REPLYAXIS queue.

```
jms:jndi:requestaxis?
 &jndiURL=iop://localhost:2810
 &jndiConnectionFactoryName=qm1
 &jndiInitialContextFactory=com.ibm.websphere.naming.WsnInitialContextFactory
 &targetService=StockQuoteAxis
 &replyToName=replyaxis;
```

- b. Add the JAR files required by the JNDI lookup.

In this configuration, the following JAR files were added to the build path to run the task using the `jndi` variant of the JMS url:

- `com.ibm.jaxws.thinclient_7.0.0.jar` from *Rational installation directory*\SDP\runtimes\base\_v7\runtimes.
- `com.ibm.ws.runtime.jar` from *Rational installation directory*\SDP\runtimes\base\_v7\plugins

For a different JNDI provider you require different JAR files.

The other JAR files in the build path are:

- 1) All the JAR files in *WebSphere MQ Installation directory*\java\lib.
- 2) All the JAR files in *Axis2-1.5.1*\lib.
- 3) Java 6.0 JRE.

- c. Run `SQA2StaticClient.java` using the same Run as... configuration as you did in the task, “Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 1249.

## Results

In both cases the reply from the service is displayed in the client console view.

## Deploying a web service client to .NET Framework 1 and 2 to use IBM MQ transport for SOAP:

Prepare a deployment directory and deployment descriptor for the client. Provide the client proxy and client class. Configure IBM MQ queues and channels, start the service and test the client.

### Before you begin

**Tip:** Develop and test the service and the client using Visual Studio. Then modify the client for IBM MQ transport for SOAP.

1. If you are deploying a service using .NET Framework 1 or 2, build the service as a library ( `.dll` ). Deploy using IBM MQ transport for SOAP.
2. Add the `Register.Extension()` call to the client.
3. Add a reference to `amqsoap.dll`, which is located in `MQ_Install\bin`.

4. Change the static `Url` property in the client proxy class constructor to the `jms:/` URI, for IBM MQ transport for SOAP.

### About this task

Deploying a web service client for .NET Framework 1 or 2 to use IBM MQ transport for SOAP requires an additional deployment step. You need to register `amqsoap.dll` with the .NET Framework. `amqsoap.dll` is automatically registered as part of installing IBM MQ transport for SOAP, but you might need to register it again.

If you use the command `amqwdeployWMQService` to generate client proxies, you can deploy the client using the directories the command generates.

### Procedure

1. Create a directory `deployDir` to hold the client deployment files.
2. Open a command window in `deployDir`.
3. Run `amqwsetcp` to set the CLASSPATH if the service is to run on Axis 1.4.
4. If necessary, run `amqwRegisterDotNet` to register `amqsoap.dll` with the .NET Framework.

### Example

The example shows the configuration and output, Figure 196 on page 1285, from an .NET Framework V2 client. The client, Figure 195, calls a web service that echoes its input parameter. The static `Url` definition, Figure 194, shows the constructor for the client proxy.

---

```
public Quote() {
 this.Url = "jms:/queue?destination=REQUESTDOTNET
 &connectionFactory=(connectQueueManager(QM1)binding(server))
 &initialContextFactory=com.ibm.mq.jms.NoJndi
 &targetService=Quote.asmx
 &replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE";
}
```

---

Figure 194. Static client proxy constructor

---

```
using System;
namespace QuoteClientProgram {
 class QuoteMain {
 static void Main(string[] args) {
 try {
 IBM.WMQSOAP.Register.Extension();
 Quote q = new Quote();
 Console.WriteLine("Response is: " + q.getQuote("ibm"));
 } catch (Exception e) {
 Console.WriteLine("Exception is: " + e);
 }
 }
 }
}
```

---

Figure 195. Client program

---

```
C:\IBM\ID\DotNet\QuoteClientProgram\QuoteClientProgram>dir /s /b
C:\IBM\ID\DotNet\QuoteClientProgram\QuoteClientProgram\QuoteClientProgram.exe
C:\IBM\ID\DotNet\QuoteClientProgram\QuoteClientProgram>quoteclientprogram
Response is: IBM
```

---

*Figure 196. Configuration and output*

### **What to do next**

1. If you are deploying the client as an IBM MQ MQI client, configure the client and server connection channel.
2. If you are deploying the client to a different queue manager to the service, you must make the destination queue available to the client. Configure the destination queue on the service queue manager as a cluster queue, or on the client queue manager as a remote queue definition.

## **Connect an Axis2 client to a JAX-WS service using W3C SOAP over JMS and WebSphere Application Server**

When you complete this task you will have called a JAX-WS web service running in WebSphere Application Server from an Axis2 client. The Axis2 client and WebSphere Application Server use the W3C candidate recommendation for the SOAP over JMS protocol running on IBM MQ. Use Eclipse Galileo and Rational Application Developer to build the web service client and web service, respectively.

### **Before you begin**

The task requires version 7 of Rational Software Development Environment and WebSphere Application Server. The task was created using the Rational Application Developer packaged with Rational Software Architect for WebSphere Software v7.5.5.1, and WebSphere Application Server Version 7.0 Test Environment v7.0.0.9 Update 1. You also require IBM MQ v7.0.1.3, or later.

The task builds on two other tasks, the Liberty profile, and “Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 1249. To complete these tasks your development environment already has Eclipse Galileo, the Liberty profile, the Eclipse plug-in for Liberty, and Axis2 1.4.1 installed.

Some of the steps are complex. The steps assume some familiarity with developing web service applications for WebSphere Application Server using Rational Application Developer. The processor and memory demands of the task are large. The task was performed in a VMWare Windows 7 SP1 virtual machine allocated 1.8GB of memory.

Install all the software before starting the task. The software takes about a day to download and a day to install, depending on bandwidth. The task takes at least half a day.

### **About this task**

The scenario for this task is that you have developed a stock quotation web service, StockQuoteAxis, using an open source tool, Eclipse Galileo. StockQuoteAxis is deployed using SOAP over HTTP running on an open source server, Liberty profile.

You want to bind the web services you deploy to a standards-based messaging transport, such as SOAP over JMS, or to web services reliable messaging, as well as to SOAP over HTTP. You want both the client and service to use standards-based interfaces. For this reason, although your future projects development team have implemented a solution using IBM MQ transport for SOAP, you have not gone into production.

The Axis2 client has removed the problem that the SOAP client for the IBM MQ transport for SOAP required a change from the HTTP client. The problem still remained that the service connected by the IBM MQ transport for SOAP is hosted by a special listener provided by IBM MQ: SimpleJavaListener.

With the W3C SOAP over JMS standard in candidate recommendation status, some vendors are providing support for W3C SOAP over JMS. The support enables you to deploy a web service to an application server and connect to the same service using a variety of connectivity protocols. The support provided by WebSphere Application Server v7 removes problem of having to host the web service separately in order to use a message-based SOAP transport. The use of a standards-based message transport interface, JMS, means you can develop solutions using tools from different vendors. You hope the web services tools in Eclipse will include SOAP over JMS bindings in the future.

Most of the steps are performed using Eclipse, or the management tools provided with the WebSphere products. The steps are described for a Windows environment. With slight modifications to some commands, you can perform the steps on other platforms.

The preliminary steps creating the HTTP web service, and connecting to it using Axis2 are listed. The client, and WSDL, from these steps are used to create the solution

## Procedure

1. Connect to the StockQuoteAxis web service using an Axis2 client and IBM MQ transport for SOAP
  - a. Use the Liberty profile
  - b. "Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse" on page 1249
  - c. "Deploying a web service client to Axis2 to use IBM MQ transport for SOAP" on page 1278
2. Connect to the StockQuoteAxis web service using an Axis2 client and the W3C candidate recommendation for SOAP over JMS.
  - a. Configure IBM MQ resources.
  - b. Configure WebSphere Application Server resources.
  - c. "Developing a JAX-WS EJB web service for W3C SOAP over JMS" on page 1236
  - d. "Deploying to an Axis2 client using W3C SOAP over JMS" on page 1281

---

## Developing REST applications with IBM MQ

### V 9.0.4

You can develop REST applications to send and receive messages. IBM MQ supports different REST APIs depending upon platform and capability.

The IBM MQ supported options you can choose from to send messages to, and receive messages from, IBM MQ using REST are the following options:

- IBM MQ messaging REST API
- IBM MQ bridge for HTTP
- IBM z/OS Connect EE
- IBM Integration Bus
- DataPower®

### V 9.0.4 IBM MQ messaging REST API

The messaging REST API comes as standard with IBM MQ from IBM MQ Version 9.0.4 and is enabled by default. You can use the messaging REST API to send and receive IBM MQ messages in plain text format.

Applications can issue an HTTP POST to send a message to IBM MQ, or an HTTP DELETE to destructively get a message from IBM MQ. Support is provided for a number of different HTTP headers which can be used to set common message properties.

The messaging REST API is fully integrated with IBM MQ security. Users must be authenticated to the mqweb server and must be a member of the MQWebUser role.

For further information, see [Messaging using the REST API](#).

## IBM MQ bridge for HTTP

The IBM MQ bridge for HTTP is a JEE application that can be installed into a suitable run time environment and provides a basic REST API on top of both the queues and topics hosted by a single MQ queue manager.

Applications can issue an HTTP POST to the bridge to send a message to IBM MQ, an HTTP GET to browse a message from IBM MQ, or an HTTP DELETE to destructively get a message from IBM MQ. Support is provided for a number of different HTTP headers which can be used to set message properties.

**Note:** From IBM MQ Version 8.0, the IBM MQ bridge for HTTP is deprecated. The IBM messaging REST API provided from IBM MQ Version 9.0.4 should be used as an alternative.

For further information, see [Developing web services with IBM MQ bridge for HTTP](#).

## IBM z/OS Connect EE

IBM z/OS Connect EE (zCEE) is a z/OS product that allows you to build REST APIs on top of existing z/OS assets such as CICS or IMS transactions, and IBM MQ queues and topics. The existing z/OS asset is hidden from the user. This allows you to REST enable your assets without changing them or any of the existing applications that use them.

With zCEE you can easily build up a REST API that will send and receive JSON data and optionally transform that data into the more traditional language structures expected by many mainframe applications. For example, COBOL copybooks.

An Eclipse based API editor can be used to build up a comprehensive RESTful API making use of query parameters and URL path segments, manipulating the JSON format as it flows through the zCEE runtime.

zCEE can be used to expose IBM MQ queues and topics as services. Two different service types are supported:

- One way services: the function provided is similar to that provided with the IBM MQ bridge for HTTP where an HTTP POST sends a message, an HTTP DELETE destructively gets a message. The main benefits over the bridge is the built in data-conversion support and that you can use the API editor to build up a more comprehensive RESTful API.
- Two way services: this provides a REST API on top of a pair of queues used by a back-end request-response style application. Callers issue an HTTP POST to the two-way service and send in JSON data. This data is put onto a request queue where it is processed by the back-end application and a response placed on the response queue. This response is retrieved and sent back to the caller as the POST response body.

zCEE is supported on IBM MQ Version 8 and later. For further information, see [IBM MQ for z/OS Service Provider for z/OS Connect](#).

## IBM Integration Bus

IBM Integration Bus is IBM's leading integration technology which can be used to connect applications and systems together regardless of the message formats and protocols that they support.

IBM Integration Bus has always supported IBM MQ and provides *HTTPInput* and *HTTPRequest* nodes which can be used to construct a RESTful interface on top of IBM MQ, and many other systems such as databases.

IBM Integration Bus can be used to do much more than provide a simple REST interface on top of IBM MQ. Its capabilities can be used to provide advanced payload manipulation, payload enrichment, and many other enhancements as part of a REST API.

For further information, see the technology sample which exposes a JSON over REST interface on top of an IBM MQ application that expects an XML payload.

## DataPower

The DataPower gateway is a single multi-channel gateway which helps provide security, control, integration and optimized access to a range of systems, including IBM MQ. It comes in both hardware and virtual form factors.

One of the services that DataPower provides is a multi-protocol gateway which can take input in one protocol and generate output in a different protocol. In particular DataPower can be configured to accept HTTP(S) data and route it to IBM MQ over a client connection, which can be used to build a REST interface on top of IBM MQ. Other DataPower services such as transformation can also be used to enhance the REST interface.

For further information, see Multi-Protocol Gateway service.

## Messaging using the REST API

▶ V 9.0.4

You can use the messaging REST API to send and receive IBM MQ messages. Information is currently sent to, and received from, the messaging REST API in plain text format.

### Before you begin

#### Note:

The messaging REST API is enabled by default. You can manually disable the messaging REST API to prevent all messaging. For more information about enabling or disabling the messaging REST API, see [Configuring the messaging REST API](#).

The messaging REST API is integrated with IBM MQ security. The caller must be authenticated to the mqweb server and must be a member of the MQWebUser role. The caller must also be authorized to access the specified queue. For more information about security for the REST API, see [IBM MQ Console and REST API security](#).

### Procedure

- “Getting started with the messaging REST API” on page 1289
- “Using the messaging REST API” on page 1291
- “Messaging REST API limitations” on page 1292
- REST API error handling
- REST API discovery
- REST API national language support



## Related information:

Messaging REST API reference

## Getting started with the messaging REST API

V 9.0.4

Before you can start the messaging REST API you must install the correct components, enable the REST API, configure security, and start the mqweb server.





### About this task

The procedure for this task focuses on getting started quickly with the messaging REST API. The steps for configuring security outline how to set up a basic user registry, but other options for configuring users and roles exist. For more information about configuring security for the messaging REST API, see IBM MQ Console and REST API security.



**Note:** You must be a privileged user to access the `mqwebuser.xml` file.

### Procedure

1. Install the IBM MQ Console and REST API component:

-  On AIX, install the `mqm.web.rte` fileset.
-  On Linux, install the `MQSeriesWeb` component. For more information about installing components and features on Linux, see Linux installation tasks.
-  On Windows, install the `Web Administration` feature. For more information about installing components and features on Windows, see Windows installation tasks.
-  On z/OS, install the `IBM MQ for z/OS Unix System Services Components` feature. For more information about installing components and features on z/OS, see z/OS installation tasks.

2. Configure users and roles before you can use the messaging REST API:

- a. Copy the `basic_registry.xml` file from the `MQ_INSTALLATION_PATH /web/mq/samp/configuration` directory.
- b. Place the sample XML file in the appropriate directory:
  -  On UNIX, Linux, and Windows: `MQ_DATA_DIRECTORY/web/installations/installationName/servers/mqweb`
  -  On z/OS: `WLP_user_directory/servers/mqweb`  
where `WLP_user_directory` is the directory that was specified when the `crtmqweb.sh` script ran to create the `mqweb` server definition.
- c. Rename the sample XML file to `mqwebuser.xml`.

**Note:** This renamed file replaces an existing file that is also used for the IBM MQ Console. Therefore, if you changed the `mqwebuser.xml` file for the IBM MQ Console, copy your changes to the new XML file before you rename it.

- d. Optional: Edit the `mqwebuser.xml` file to add users and groups. Assign those users and groups to the `MQWebUser` role to be authorized to use the messaging REST API. The `MQWebAdmin` and `MQWebAdminR0` roles are not applicable for the messaging REST API. You can also change the passwords for the users that are defined by default, and encode the new passwords. Ensure that the users are authorized to access the specified queue. For more information, see [Configuring users and roles](#).
3. Enable remote connections to the `mqweb` server by using the `setmqweb` command:
- ```
setmqweb properties -k httpHost -v hostname
```

where *hostname* specifies the IP address, domain name server (DNS) host name with domain name suffix, or the DNS host name of the server where IBM MQ is installed. Use an asterisk, *, to specify all available network interfaces.

Attention:  

Before issuing either the **setmqweb** or **dspmqweb** commands on z/OS, you must set the WLP_USER_DIR environment variable, so that the variable points to your mqweb server configuration.

To do this, issue the following command:



```
export WLP_USER_DIR=WLP_user_directory
```

where *WLP_user_directory* is the name of the directory that is passed to `crtmqweb.sh`. For example:

```
export WLP_USER_DIR=/var/mqm/web/installation1
```

For more information, see [Create the Liberty server definition](#).

4. Start the mqweb server that supports the REST API:

-  As a privileged user, enter the following command on the command line:
`strmqweb`
-  On z/OS, start the procedure that you created in Task 29: [Create a procedure for the IBM WLP server](#).

What to do next

1. Choose how users of the messaging REST API authenticate with the mqweb server. You do not have to use the same method for all users. The following options are available:
 - Let users authenticate by using HTTP basic authentication. In this case, a user name and password are encoded, but not encrypted, and sent with each REST API request to authenticate and authorize the user for that request. In order for this authentication to be secure, you must use a secure connection. That is, you must use HTTPS. For more information, see [Using HTTP basic authentication with the REST API](#).
 - Let users authenticate by using token authentication. In this case, a user provides a user ID and password to the REST API login resource with the HTTP POST method. An LTPA token is generated that enables the user to remain logged in and authorized for a set amount of time. In order for this authentication to be secure, you must use a secure connection. That is, you must use HTTPS. For more information, see [Using token based authentication with the REST API](#).
 - Let users authenticate by using client certificates. In this case, the user does not use a user ID or password to log in to the messaging REST API, but uses the client certificate instead. For more information, see [Using client certificate authentication with the REST API](#).
2. Configure REST API settings, including enabling HTTP connections, and changing the port number. For more information, see [Configuring the IBM MQ Console and REST API](#).
3. Optionally, configure Cross Origin Resource Sharing for the REST API. By default, you cannot access the REST API from web resources that are not hosted on the same domain as the REST API. That is, cross-origin requests are not enabled. You can configure Cross Origin Resource Sharing (CORS) to allow cross-origin requests from specified URLs. For more information, see [Configuring CORS for the REST API](#).
4. Use the REST API. For more information, see [“Using the messaging REST API” on page 1291](#), and the [Messaging REST API reference](#).

Note: You can stop the mqweb server at any time by using the **endmqweb** command. However, if the mqweb server is not running, you cannot use the REST API or IBM MQ Console.

Using the messaging REST API

V 9.0.4

When you use the messaging REST API, you invoke HTTP methods on URLs to send and receive IBM MQ messages. The HTTP method, for example POST, represents the type of action to be performed on the object that is represented by the URL. Further information about the action might be encoded in query parameters. Information about the result of performing the action might be returned as the body of the HTTP response.

Before you begin

Consider these things before you use the messaging REST API:

- You must authenticate with the mqweb server in order to use the messaging REST API. You can authenticate by using HTTP basic authentication, client certificate authentication, or token based authentication. For more information about how to use these authentication methods, see IBM MQ Console and REST API security.
- The REST API is case-sensitive. For example, an HTTP POST on the following URL results in an error if the queue manager is called qmgr1.
`/ibmmq/rest/v1/messaging/qmgr/QMGR1/queue/Q1/message`
- Not all of the characters that can be used in IBM MQ object names can be directly encoded in a URL. To encode these characters correctly, you must use the appropriate URL encoding:
 - A forward slash, /, must be encoded as %2F.
 - A percent sign, %, must be encoded as %25.

About this task

When you use the REST API to perform a messaging action on an IBM MQ queue object, you first need to construct a URL to represent that object. Each URL starts with a prefix, which describes which host name and port to send the request to. The rest of the URL describes a particular object, or route to that object, known as a resource.

The messaging action that is to be performed on the resource defines whether the URL needs query parameters or not. It also defines the HTTP method that is used, and whether additional information is sent to the URL, or returned from it. The additional information might form part of the HTTP request, or be returned as part of the HTTP response.

After you construct the URL, you can send the HTTP request to IBM MQ. You can send the request by using the HTTP implementation that is built into the programming language of your choice. You can also send the request by using command line tools such as cURL, or a web browser, or web browser add-on.

Important: You must, as a minimum, carry out steps 1a and 1b on page 1292.

Procedure

1. Construct the URL:

- a. Start with the following prefix URL:

```
https://host:port/ibmmq/rest/v1/messaging
```

host

Specifies the host name or IP address that the messaging REST API is available on.

The default value is localhost.

port

Specifies the HTTPS port number that the messaging REST API uses.

The default value is 9443.

If you enable HTTP connections, you can use HTTP instead of HTTPS. For more information about enabling HTTP, see [Configuring HTTP and HTTPS ports](#).

For more information about how to determine the prefix URL, see [Determining the REST API URL](#).

- b. Add the queue and associated queue manager resources to use for messaging to the URL path. In the messaging reference, the variable segments can be identified in the URL by the braces that surround it { }. For further information, see `/messaging/qmgr/{qmgrName}/queue/{queueName}/message`.

For example, to interact with queue *Q1* associated with queue manager *QM1*, add `/qmgr` and `/queue` to the prefix URL to create the following URL:

```
https://localhost:9443/ibmmq/rest/v1/messaging/qmgr/QM1/queue/Q1/message
```

- c. Optional: Add an optional query parameter to the URL.

Add a question mark, `?`, query parameter, equal sign `=`, and a value to the URL.

For example, to wait for a maximum of 30 seconds for the next message to become available, create the following URL:

```
https://localhost:9443/ibmmq/rest/v1/messaging/qmgr/QM1/queue/Q1/message?wait=30000
```

- d. Optional: Add further optional query parameters to the URL.

Add an ampersand, `&`, to the URL, and then repeat step 1c.

2. Invoke the relevant HTTP method on the URL. Specify any optional message payload, and provide the appropriate security credentials to authenticate. For example:
 - Use the HTTP/REST implementation of your chosen programming language.
 - Use a tool such as a REST client browser add-on or cURL.

Messaging REST API limitations

▼ 9.0.4

Before you use the messaging REST API, consider the following limitations:

- The API currently does not support publish/subscribe messaging. The messaging REST API supports only synchronous point-to-point messaging.
- The API currently does not support browsing from queues. Messages are destructively received from the IBM MQ queue using the HTTP DELETE method. For further information, see [DELETE](#).
- When sending a message only UTF-8 text based content can be used. Messages are put as IBM MQ MQSTR formatted messages. For further information, see [POST](#).

When receiving a message only IBM MQ MQSTR formatted messages are supported. Subsequently, all messages are received under sync-point and any unhandled messages are left on the queue.

The IBM MQ queue can be configured to move these poison messages to an alternate destination. For further information, see [Handling poison messages in IBM MQ classes for JMS](#).

- If you use Advanced Message Security (AMS) with the messaging REST API, note that all messages are encrypted by using the context of the mqweb server, not the context of the user that posts the message.

Developing web services with IBM MQ bridge for HTTP

V 9.0.4

With IBM MQ bridge for HTTP, client applications can exchange messages with IBM MQ without the need to install an IBM MQ MQI client. You can call IBM MQ from any platform or language with HTTP capabilities.

About this task

Note: From IBM MQ Version 8.0, the IBM MQ bridge for HTTP is deprecated. The IBM messaging REST API provided from IBM MQ Version 9.0.4 should be used as an alternative.

Introduction to IBM MQ bridge for HTTP

V 9.0.4

The IBM MQ bridge for HTTP is a Java, Enterprise Environment (JEE) Web application. HTTP clients can send **POST**, **GET**, and **DELETE** requests to it to put, browse and delete messages from IBM MQ queues. The IBM MQ bridge for HTTP is not suitable for use with messages, if assured delivery is required.

Benefits

With the IBM MQ bridge for HTTP you can send and receive IBM MQ messages using HTTP from a wide variety of environments:

- Environments that support HTTP, but not IBM MQ.
- Environments that have insufficient storage space to install an IBM MQ MQI client.
- Environments that are too numerous to install the IBM MQ MQI client on each system that requires access to IBM MQ.
- Web-based applications from which you want to send or receive messages without coding your own bridge to IBM MQ.
- Web-based applications that you want to enhance, using asynchronous techniques such as AJAX. IBM MQ bridge for HTTP makes IBM MQ queues and topics available using Representation State Transfer (REST) over HTTP.

HTTP support can be used with both point-to-point and publish/subscribe messaging topologies.

How does HTTP support work?




Figure 197. IBM MQ bridge for HTTP

The IBM MQ bridge for HTTP Web application receives HTTP requests from one or more clients. It interacts with IBM MQ on their behalf, and returns HTTP responses to them.

The IBM MQ bridge for HTTP is a JEE servlet that is connected to IBM MQ using a resource adapter. The HTTP servlet handles three different types of HTTP requests: **POST**, **GET**, and **DELETE**.

Table 186. IBM MQ bridge for HTTP verbs

| HTTP Request | Result |
|--------------|--|
| POST | Puts a message on a queue or topic. |
| GET | Browses the first message on a queue. In line with the HTTP protocol, GET does not delete the message from the queue. Do not use GET with publish/subscribe messaging. |
| DELETE | Gets and deletes a message from a queue or topic. |

HTTP POST example

HTTP **POST** puts a message to a queue, or a publication to a topic. The **HTTPPOST** Java sample is an example an HTTP **POST** request of a message to a queue. Instead of using Java, you could create an HTTP **POST** request using a browser form, or an AJAX toolkit instead.

The following figure shows an HTTP request to put a message on a queue called myQueue. This request contains the HTTP header `x-msg-correlID` to set the correlation ID of the IBM MQ message.

```
POST /msg/queue/myQueue/ HTTP/1.1
Host: www.example.org
Content-Type: text/plain
x-msg-correlID: 1234567890
Content-Length: 50

Here is my message body that is posted on the queue.
```

Figure 198. Example of an HTTP **POST** request to a queue

The following figure shows the response sent back to the client. There is no response content.

```
HTTP/1.1 200 OK
Date: Wed, 2 Jan 2007 22:38:34 GMT
Server: Apache-Coyote/1.1 WMQ-HTTP/1.1 JEE-Bridge/1.1
Content-Length: 0
```

Figure 199. Example of an HTTP **POST** response

HTTP DELETE example

HTTP **DELETE** gets a message from a queue and deletes the message, or retrieves and deletes a publication. The **HTTPDELETE** Java sample is an example an HTTP **DELETE** request reading a message from a queue. Instead of using Java, you could create an HTTP **DELETE** request using a browser form, or an AJAX toolkit instead.

The following figure shows an HTTP request to delete the next message on queue called myQueue. In response, the message body is returned to the client. In IBM MQ terms, HTTP **DELETE** is a destructive get.

The request contains the HTTP request header `x-msg-wait`, which instructs IBM MQ bridge for HTTP how long to wait for a message to arrive on the queue. The request also contains the `x-msg-require-headers` request header, which specifies that the client is to receive the message correlation ID in the response.

```
DELETE /msg/queue/myQueue/ HTTP/1.1
Host: www.example.org
x-msg-wait: 10
x-msg-require-headers: correlID
```

Figure 200. Example of an HTTP **DELETE** request

The following figure shows the response returned to the client. The correlation ID is returned to the client, as requested in `x-msg-require-headers` of the request.

```
HTTP/1.1 200 OK
Date: Wed, 2 Jan 2007 22:38:34 GMT
Server: Apache-Coyote/1.1 WMQ-HTTP/1.1 JEE-Bridge/1.1
Content-Length: 50
Content-Type: text/plain; charset=utf-8
x-msg-correlId: 1234567890

Here is my message body that is retrieved from the queue.
```

Figure 201. Example of an HTTP **DELETE** response

HTTP **GET** example

HTTP **GET** gets a message from a queue. The message remains on the queue. In IBM MQ terms, HTTP **GET** is a browse request. You could create an HTTP **GET** request using a Java client, a browser form, or an AJAX toolkit.

The following figure shows an HTTP request to browse the next message on queue called `myQueue`.

The request contains the HTTP request header `x-msg-wait`, which instructs IBM MQ bridge for HTTP how long to wait for a message to arrive on the queue. The request also contains the `x-msg-require-headers` request header, which specifies that the client is to receive the message correlation ID in the response.

```
GET /msg/queue/myQueue/ HTTP/1.1
Host: www.example.org
x-msg-wait: 10
x-msg-require-headers: correlID
```

Figure 202. Example of an HTTP **GET** request

The following figure shows the response returned to the client. The correlation ID is returned to the client, as requested in `x-msg-require-headers` of the request.

```
HTTP/1.1 200 OK
Date: Wed, 2 Jan 2007 22:38:34 GMT
Server: Apache-Coyote/1.1 WMQ-HTTP/1.1 JEE-Bridge/1.1
Content-Length: 50
Content-Type: text/plain; charset=utf-8
x-msg-correlId: 1234567890

Here is my message body that appears on the queue.
```

Figure 203. Example of an HTTP **GET** response

Installing, configuring, and verifying IBM MQ bridge for HTTP

V 9.0.4

Obtain IBM MQ bridge for HTTP by installing Java Messaging and Web Services from either the IBM MQ MQI client or server installation materials. Deploy IBM MQ bridge for HTTP to a suitable application server.


Before you begin

Check the prerequisite products at System Requirements for IBM MQ. The installation process does not check for the presence and availability of the prerequisite software for running IBM MQ bridge for HTTP. You must verify that the prerequisites are installed.

IBM MQ bridge for HTTP is a Java EE 4 application. For information about supported application servers, see System Requirements for IBM MQ.

About this task

IBM MQ bridge for HTTP is supplied as a .war file, WMQHTTP.war.

- On UNIX and Linux:
 - WMQHTTP.war is included as part of the Java Messaging and Web Services installation option. The option is available in both the client and server installation materials.
 - WMQHTTP.war is installed to *mqmtop/java/http/WMQHTTP.war*. *mqmtop* is the directory where IBM MQ is installed.
 - WMQHTTP.samples is installed to *mqmtop/java/http/samples*. *mqmtop* is the directory where IBM MQ is installed.
-  On z/OS:
 - WMQHTTP.war is included as part of the IBM MQ z/OS UNIX System Services Components feature.
 - WMQHTTP.war is installed to *PathPrefix/usr/lpp/mqm/V7R0M0/HTTPBridge/*, where *PathPrefix* is an optional customer defined prefix.

Carry out the following installation steps to install IBM MQ bridge for HTTP, deploy, and configure it, and verify the configuration. The details of the configuration steps vary on different application servers. Use “Deploying and verifying IBM MQ bridge for HTTP on WebSphere Application Server V6.1.0.9” on page 1297 as a template for the steps to follow on your application server.

Procedure

1. Obtain WMQHTTP.war by installing either the IBM MQ MQI client or server.
2. Copy WMQHTTP.war to a server from which it can be deployed to an application server.
3. Deploy WMQHTTP.war to an application server.
4. If necessary, install IBM MQ as a resource adapter on your application server. Find out if IBM MQ is already configured as a messaging provider on your application server. Use the administration or management tool supplied with your application server, to look for IBM MQ. IBM MQ might be found under the following path, **Resources > JMS > Messaging providers**.
5. Configure a connection factory on the application server to connect to a queue manager that uses the IBM MQ MQI client transport¹².
6. Configure the WMQHTTP.war Web application on the application server to use the connection factory
7. Verify the configuration.

12. Initially, at least, configure the client transport. Some application servers can connect to IBM MQ by using direct, or bindings mode connections.

- a. Set up the queue manager named in the connection factory and a local queue.
- b. Place a message on the local queue.
- c. Create the server-connection channel named in the connection factory, with authority to read and write to the local queue.
- d. Start the queue manager and the listener.
- e. Start the application server and WMQHTTP.war, if they are not already running.
- f. Open a browser and type `http://hostname: web port/Context root/msg/queue/local queue`

Results

The browser window displays the message you placed on the local queue.

What to do next

1. Try the example, “Deploying and verifying IBM MQ bridge for HTTP on WebSphere Application Server V6.1.0.9.”
2. Run the sample HTTP Java applications.

Deploying and verifying IBM MQ bridge for HTTP on WebSphere Application Server V6.1.0.9:



Use the following example to prepare a deployment of IBM MQ bridge for HTTP to run the sample HTTP Java programs. The deployment is on WebSphere Application Server V6.1.0.9.

Before you begin

1. Follow the instructions in “Installing, configuring, and verifying IBM MQ bridge for HTTP” on page 1296, to copy WMQHTTP.war onto a server accessible to your installation of WebSphere Application Server.
2. Configure a queue manager, and a queue, to use to test the configuration:
 - In the example, the queue manager is configured as using the values in Table 187:

Table 187. Queue manager configuration

| Object | Value |
|---------------------------|--|
| Host name | itso-01 |
| Queue manager | QM1 |
| Local queue | HTTPTESTQ |
| Server connection channel | MYSVRCON. Configure an MCA user ID with sufficient authority to read and write to HTTPTESTQ. |
| Listener port | 1414 |

3. Start the queue manager and the listener
4. Place a test message onto HTTPTESTQ. For example:
 - a. Start IBM MQ Explorer.
 - b. In the list of local queues for QM1, right-click **HTTPTESTQ** > **Put test message** > **type First Message** > **Put message** > **Close**
5. Start the application server and sign on to the Integrated Solutions Console.

About this task

The example shows the steps to take if you are running WebSphere Application Server V6.1.0.9 as your application server. If you are running a different version of WebSphere Application Server, or running a

different application server, the steps are different. WebSphere Application Server V6.1.0.9 is pre-configured with IBM MQ installed as a message provider, using the IBM MQ MQI client libraries. If IBM MQ is not pre-configured as a messaging provider, or if you want to use IBM MQ server bindings, you need to install and configure the IBM MQ resource adapter for JEE into your application server.

Follow the instructions to deploy IBM MQ bridge for HTTP onto WebSphere Application Server V6.1.0.9, and verify the deployment using a browser:

Procedure

1. In the navigation pane, click **Resources > JMS providers > IBM MQ messaging provider**.
You can configure at either Node, Cell, or Server level, depending on your WebSphere Application Server deployment. The example uses Server level deployment.
2. Under Additional properties, click **Connection factories > New**.
3. In the JMS providers form, provide the information in Table 188, or alternatives of your choosing, click **Apply > Save**.

Table 188. Set or modify the following fields

| Field | Value |
|----------------|---------------------------------|
| Name | WMQHTTPBridge |
| JNDI Name | jms/WMQHTTPJCAConnectionFactory |
| Queue manager | QM1 |
| Host | itso-01 |
| Port | 1414 |
| Channel | MYSVRCON |
| Transport type | CLIENT |

4. In the navigation pane, click **Applications > Install New Application**.
5. Insert the path to WMQHTTP.war into the form, and provide a Context root, click **Next**.
 - a. The Context root is optional. mq is the default Context root for the sample HTTP applications.
 - b. The Context root forms part of the URI identifying IBM MQ bridge for HTTP. You can omit the Context root, or change it later.
6. On the Select installation options page of the installation wizard, you do not have to change any of the defaults, click **Next**.
7. On the Map modules to servers page, select a Cluster or Server, check the Select box, click **Apply > Next**.
8. On the Map resource references to resources page, in the javax.jms.ConnectionFactory form, click **Browse...** on the IBM MQ bridge for HTTP row.
9. On the Enterprise Applications > Available resources page, select **WMQHTTPBridge**, click **Apply**.
10. Back in the javax.jms.ConnectionFactory form, select the authentication method.
 - a. For the example, choose **None**, click **Apply**. The other options require additional configuration.
11. Check the **Select** check box for IBM MQ bridge for HTTP, click **Next > Next > Finish > Save**
12. In the navigation pane, click **Applications > Enterprise Applications**.
13. Check the selection box for WMQHTTP.war, click **Start**.
14. Open a browser window. Type `http://itso-01:9080/mq/msg/queue/HTTPTESTQ`, using the appropriate host name and port.

Results

The browser window displays First Message, if the configuration is successful.

What to do next

Run the sample HTTP Java applications.

Publish/subscribe using the IBM MQ bridge for HTTP

► V 9.0.4

IBM MQ bridge for HTTP uses the IBM MQ classes for JMS publish/subscribe interface. HTTP **POST** creates a publication. HTTP **DELETE** creates a non-durable managed subscription. You must configure publish/subscribe for JMS before using the topic URI.

Publish/subscribe is fully integrated into IBM MQ in version 7. Before version 7, a separate publish/subscribe broker handled publications and subscriptions. It is called “queued” publish/subscribe, to distinguish it from the fully integrated publish/subscribe in version 7. Version 7 emulates queued publish/subscribe using integrated publish/subscribe. The emulation enables existing queued publish/subscribe applications to coexist with integrated applications running on the same queue manager. Queued publish/subscribe applications can also interoperate with integrated applications, sharing the same topics. In Version 6, the broker was shipped with IBM MQ; before Version 6 it was available as a SupportPack.

Configuration

The IBM MQ bridge for HTTP uses the JMS interface to publish and subscribe. In version 7, you can control whether the IBM MQ classes for JMS use queued or integrated publish/subscribe, using the PROVIDERVERSION JMS property.

An additional consideration is that you can use either IBM MQ MQI client libraries with IBM MQ bridge for HTTP, or server libraries. Version 6 client libraries only support queued publish/subscribe, whereas version 7 libraries support both queued and integrated publish/subscribe. Most Web or application servers that use IBM MQ as a messaging provider do so using client libraries. In order to use integrated publish/subscribe, both the IBM MQ MQI client and server libraries must be at least at version 7. If either is running an earlier version of WebSphere than 7, then you must configure queued publish/subscribe; see Table 189. Check what libraries are installed or configured with the Web server or application server you are using.

Table 189. Publish/subscribe configuration modes

| | Client V6 or earlier | Client V7 or later |
|----------------------|--|--|
| Server V6 or earlier | 1. Run the <code>\java\bin\MQJMS_PSQ.mqsc</code> script | Not supported |
| Server V7 or later | 1. Run the <code>\java\bin\MQJMS_PSQ.mqsc</code> script
2. Set the queue manager to <code>PSMODE=ENABLED</code> | 1. If <code>PROVIDERVERSION = 7</code>
a. Set the queue manager to <code>PSMODE=ENABLED</code> or <code>PSMODE=COMPAT</code>
2. If <code>PROVIDERVERSION = 6</code>
a. Set the queue manager to <code>PSMODE=ENABLED</code> |

Publish

Send an HTTP **POST** request with the URI:

```
http://hostname:port/context_root/msg/topic/topicString
```

The message contents are published using the topic string *topicString*.

Subscribe

Send an HTTP **DELETE** request with the URI:

`http://hostname:port/context_root/msg/topic/topicString`

IBM MQ bridge for HTTP creates a managed non-durable subscription to the topic string *topicString*. The subscription is deleted as soon as a publication is returned, or until the wait-interval set by the custom entity-header, *x-msg-wait*, expires.

Running the IBM MQ bridge for HTTP samples

V 9.0.4






The IBM MQ bridge for HTTP samples are available for use on only the Windows operating system. The samples show you how to submit HTTP **POST** and HTTP **DELETE** commands to IBM MQ bridge for HTTP from Java programs.

Before you begin

Verify your IBM MQ bridge for HTTP installation by running step 7 on page 1296 in “Installing, configuring, and verifying IBM MQ bridge for HTTP” on page 1296.

The HTTP samples are installed to the directories shown in Table 190. In each case, source code is installed to the */src* subdirectory.

Table 190. Location of HTTP samples

| Platform | Location |
|---|---|
|  Windows | <code>MQ_INSTALLATION_PATH/tools/http/samples</code> |
|   z/OS | <code>PathPrefix/usr/lpp/mqm/V7R0M0/http/samples</code> |
|   IBM i | <code>MQ_INSTALLATION_PATH/java/samples/http</code> |
| All other platforms | <code>MQ_INSTALLATION_PATH/samp/http</code> |


MQ_INSTALLATION_PATH represents the directory where IBM MQ is installed.

About this task

The samples simulate the IBM MQ AMQSPUT and AMQSGET sample applications. They illustrate the following functions in a point-to-point messaging environment:

- **HTTPPOST** - Sends HTTP **POST** requests in a Java application to put messages to an IBM MQ queue, using the IBM MQ bridge for HTTP and handles the responses.
- **HTTPDELETE** - Sends HTTP **DELETE** requests in a Java application to get messages from an IBM MQ queue, using the IBM MQ bridge for HTTP and handles the responses containing the IBM MQ message.

Parameters for HTTPPOST and HTTPDELETE



To run the **HTTPPOST** sample, complete the following steps:

Procedure

1. In a command window, navigate to the HTTP samples directory.
2. Run the **HTTPPOST** sample.

```
java -classpath . HTTPPOST [parameters]
```

When the **HTTPPOST** sample starts, the following output is displayed:

```
HTTP POST Sample start
Target server is ' hostName '
Target queue is ' queueName '
Target context-root is ' contextRoot '
```

3. In the command prompt, type the text that you want to form your message body.
4. Press Enter to post the message to the IBM MQ queue.
 - a. If you want to send another message, enter some more text. The text forms the body of a second IBM MQ message.
 - b. Press Enter to post the message to the IBM MQ queue.
5. Press Enter twice to end **HTTPPOST**. The following output is displayed:

```
HTTP POST Sample end
```

What to do next

The **HTTPDELETE** sample performs a destructive get of all the messages you placed on the IBM MQ queue.

Run the **HTTPDELETE** sample by completing the following steps:

1. In a command window, navigate to `MQ_INSTALLATION_PATH/tools/samples`. `MQ_INSTALLATION_PATH` represents the directory where IBM MQ is installed.
2. Run the **HTTPDELETE** sample.

```
java -classpath . HTTPDELETE [parameters]
```

When the **HTTPDELETE** sample starts, the following output is displayed:

```
HTTP DELETE Sample start
Target server is ' host:port '
Target queue is ' your queue name '
Target context-root is ' your context-root '
message
message
...
```

Security considerations for IBM MQ bridge for HTTP

V 9.0.4

Standard web security considerations apply to authenticating a web browser client. Authorization to IBM MQ resources is at the level of the user running the IBM MQ bridge for HTTP servlet, and not the individual web browser client. Standard IBM MQ security consideration apply to IBM MQ.

Data flowing from a web browser to an IBM MQ application using IBM MQ bridge for HTTP, and back, takes a three steps:

Client connection

From the browser to the IBM MQ bridge for HTTP over a TCP/IP connection using HTTP.

Resource adapter connection to IBM MQ

The connection is from the IBM MQ bridge for HTTP to an IBM MQ queue manager. The connection is either a client connection, over TCP/IP, or a local IBM MQ bindings connection. Once the connection is made, the HTTP request is placed on a standard local queue or a transmission queue.

From the IBM MQ local queue over one or more channels, to the target queue.

Apply standard techniques for securing queues, topics, queue managers, and channels.

The reply takes the steps in reverse.

Client connection

Secure connections between HTTP clients and the application server using the web container. Use standard HTTP server techniques, such as using HTTPS. Refer to the documentation for your application server for information.

Resource adapter connection to IBM MQ

The connection between the resource adapter and queue manager is authorized using only a single user ID. Assign a single user ID to identify requests from the IBM MQ bridge for HTTP. The user ID must have restricted IBM MQ authorizations only to the resources external users must have access. You must authenticate the actual client separately, and establish trust for successive interactions with the client, using standard techniques for web security.

Secure the connection between the resource adapter and the queue manager using the single user ID. Restrict the authorities the user ID has to no more than needed to read and write messages to queues and topics. The IBM MQ bridge for HTTP is a point of attack between the internet and your intranet.

How you secure the connection between your resource adapter and IBM MQ is dependent on your specific resource adapter. Refer to the documentation for the resource adapter.

Index

Special characters

- .NET
 - support for 1
- .NET Monitor 675

Numerics

- 64 bit C applications, building on
 - z/OS 421

A

- About IBM MQ automation classes for
 - ActiveX classes 689
- AccessGetMessageOptions method 693
- accessing queues
 - classes for .NET 654
- accessing queues, topics, and processes in
 - Java 1050
- accessing topics
 - classes for .NET 654
- AccessPutMessageOptions method 693
- AccessQueue method 706
- AccessQueueManager method 693
- AccountingToken field 40
- AccountingToken property 734, 770
- AccountingTokenHex property 734, 771
- AccessQueue method 694
- acknowledgment mode
 - JMS 944
- ActiveX
 - support for 1
- adapter
 - batch 261
 - IMS 262
- Add method 785
 - Selector parameter 785
 - Value parameter 785
- AddDistributionList method 707
- AddDistributionListItem method 767
- AddInquiry method 786
 - Inquiry parameter 786
- Address space models
 - HP-UX on IA64 (IPF) 398
- administered objects 871
 - retrieving from JNDI 931
- advantages
 - IBM MQ classes for JMS 824
- AIX operating system
 - performance of nonpersistent messages 151
- alias queue
 - resolving queue name 104
- alias queue manager definition 102
- alternate PCB, IMS bridge 68, 281
- AlternateUserId field 108
- AlternateUserId property 696, 711, 764
- alternative user authority 108
- amasps0 sample program 502

- AMQ0ECHA sample program 491
- amq0gbr0 sample program 459, 478
- amq0get0 sample program 459, 492
- AMQ0GET4 sample program 492
- amq0put0 sample program 459, 503
- amq0req0 sample program 459, 512
- AMQ0REQ4 sample program 512
- AMQ0SETA sample program 518
- amqiech2 sample program 465, 491
- amqiechx sample program 459, 491
- amqiinq2 sample program 465, 496
- amqiinqx sample program 459, 496
- amqiset2 sample program 465, 518
- amqisetx sample program 459, 518
- amqltmc0 trigger monitor 250
- amqmec2 sample program 465, 491
- amqmecx sample program 459, 491
- amqminq2 sample program 465, 496
- amqminqx sample program 459, 496, 497
- amqmset2 sample program 465, 518
- amqmsetx sample program 459, 518
- AMQP
 - messages 1172
- AMQP CorrelationID header field 1172
- amqrgrma sample program 505
- AMQSAPT sample program 477
- amqsapt0 sample program 477
- amqsaxe0 sample program 459, 475
- amqsbcg sample program 480
- amqsbcg0 sample program 459, 480
- amqsbcgc sample program 480
- amqsbcf 476
- amqscic0 sample transaction 459, 481
- amqscic21 sample transaction 459
- AMQSCLM sample program 540, 547
- amqsclma.c sample program 540, 547
- amqscnxb.vbp sample program 450
- amqsdlq sample program 459
- amqsech sample program 491
- amqsecha sample program 459, 491
- AMQSECHA sample program 491
- amqsechc sample program 491
- AMQSERV4 sample program 248, 522
- amqsgbr sample program 478
- amqsgbr0 sample program 459, 478
- AMQSGBR4 sample program 478
- amqsgbrc sample program 478
- amqsget0 sample program 459, 492
- AMQSGET4 sample program 492
- amqsgetc sample program 459, 492
- amqsinqa sample program 496
- amqsinqc sample program 496
- amqsiqua sample program 497
- amqsprma sample program 505
- amqsptl0 sample program 490
- amqsput0 sample program 459, 503
- AMQSPUT4 sample program 503
- amqsputc sample program 459, 503
- amqsreq sample program 512
- amqsreq0 sample program 459, 512

- AMQSREQ4 sample program 512
- amqsreqc sample program 512
- amqsset sample program 518
- amqsseta sample program 518
- AMQSSETA sample program 518
- amqssetc sample program 518
- amqsspin sample program 459
- AMQSSSLC 520
- amqstrg sample program 522
- amqstrg0 sample program 459, 522
- AMQSTRG4 sample program 248, 522
- amqstrgc sample program 522
- amqstxgx sample program 459, 524
- amqstxpx sample program 459, 524
- amqstxsx sample program 459, 524
- amqsvfc0 sample program 465, 482
- AMQSVFC4 sample program 468, 482
- amqsvfcx sample program 482
- amqswlm0 sample program 459
- amqsxa4x sample transaction 459
- amqsxab0.sqb sample 483
- amqsxab0.sqc sample 483
- amqsxaex sample transaction 459
- amqsxaf0.sqb sample 483
- amqsxaf0.sqc sample 483
- amqsxag0.c sample 483
- amqsxag0.cb1 sample 483
- amqsxas0.sqb sample 483
- amqsxas0.sqc sample 483
- amqsxrma sample program 505
- API (Application Programming Interface)
 - calls 76
 - dealing with failure of a call 436
- API exits
 - amqsaxe0 sample program 475
- API-crossing exit for z/OS
 - CSQCAPX 265
 - introducing 263
 - preparing 266
 - sample 265
 - writing your own 263
- application configuration file
 - .NET 650
- application data 9
- application design, performance considerations 53
- Application log 645
- application programming
 - application takeover 267
 - connection tag 267
 - dynamic queues 268
 - index queues 268
 - migrating applications to use shared queues 269
 - persistent messages 268
 - queue-sharing groups 266
 - serialized applications 267
- application programs
 - design considerations 48
 - MQI local administration, support for 5

- application programs (*continued*)
 - persistent messages, effect on performance 48
 - searching for messages, effect on performance 48
 - threads, application design 48
 - Application Server Facilities 1002
 - handling poison messages in ASF 1006
 - ApplicationIdData property 734
 - ApplicationOriginData property 735
 - applications
 - building
 - clients 291
 - connected to multiple queue managers 294
 - connected to multiple servers 294
 - connection to server 294
 - in different environments 294
 - JMS, writing 871
 - running 1086
 - WebSphere
 - MQ 42
 - applications that access non-ActiveX applications 682
 - applications, writing 283
 - AppIdentityData field 40
 - AppOriginData field 40
 - AppType 306
 - ASF (Application Server Facilities) 1002
 - assembler language
 - preparing your program to run 420
 - support for 452
 - using the MQI 455
 - asynchronous message delivery 950
 - asynchronous put
 - JMS 993
 - attributes 558
 - DefInputOpenOption 107
 - DefPriority 17
 - HardenGetBackout 39, 438
 - IndexType 152
 - inquiring about 218
 - MaxMsgLength 116, 133
 - MaxPriority 17
 - MsgDeliverySequence 134
 - ProcessName 240
 - selectors 219
 - setting 218
 - Shareability 107
 - TrigData 240
 - TriggerControl 245
 - TriggerData 235
 - TriggerDepth 246
 - TriggerMsgPriority 245
 - TriggerType 245
 - attributes of objects 803
 - audit trail using message context 40
 - authority checking
 - alternative user authority on MQOPEN 108
 - by MQOPEN 100
 - AuthorityEvent property 696
 - authorization
 - refreshing the OAM after changing a user's 326
 - authorization service
 - component 330
 - stanza 331
 - stanza, Windows 325
 - user interface 319, 331
 - AUTO_ACKNOWLEDGE 944
 - auto-definition exit program 362
 - automatic client reconnection 297
 - Automatic client reconnection 662
 - automatically starting an application
 - an example 514
 - how triggering works 234
- B**
- backing out changes 39, 222
 - Backout method 707
 - backout, skipping 161
 - BackoutCount field 39
 - BackoutCount property 735
 - BackoutRequeueName property 712
 - BackoutThreshold property 712
 - BaseQueueName property 712
 - batch for IBM MQ for z/OS
 - adapter 261
 - building an application 421
 - calling the stub dynamically 427
 - restrictions 87
 - support for 259
 - Begin method 707
 - BeginOptions property 697
 - behavior in different environments 1092
 - benefits
 - IBM MQ classes for JMS 824
 - binary strings 804, 806
 - binding
 - FASTPATH 94, 99
 - STANDARD 94, 99
 - bindings
 - connection 851, 1022
 - connection, programming
 - Java 1044
 - verifying 1040
 - bindings mode 636
 - bindings mode connections 1021
 - body of a message
 - types 949
 - body, message 875, 1172
 - BROKERCCDSUBQ object
 - property 1004
 - BROKERCCSUBQ object property 1004
 - browse cursor 106
 - browse with mark 171
 - browsing (sample for IBM MQ for z/OS) 590
 - browsing messages 165
 - browsing messages in logical order 168
 - bufferPointer method 811
 - buffers, message 807
 - building a connection 942
 - building applications
 - z/OS 821
 - z/OS UNIX System Services 823
 - building client applications 291
 - building your application
 - batch with IBM MQ for z/OS 421
 - CICS and IBM MQ for z/OS 425
 - building your application (*continued*)
 - IMS 426
 - on AIX 382
 - on HP-UX 393
 - on IBM i 404
 - on Linux 398
 - 31-bit 400
 - 32-bit 401
 - 64-bit 402
 - on Solaris 407
 - on Windows systems 413
 - on z/OS 420
 - UNIX System Services 427
 - XPLINK with IBM MQ for z/OS 423
 - bundles, OSGi 862
 - bundles, OSGi - IBM MQ classes for
 - Java 1031
 - bytes message 875, 1172
- C**
- C applications
 - linking 292
 - C Extended Transaction client 639
 - C language
 - support for 448
 - C Set++ 816
 - C, using from C++ 805
 - C++
 - support for 1
 - unsupported functions 807
 - C++ applications
 - linking 293
 - C++ language considerations 803
 - call interface 76
 - calling dynamically with IBM MQ for z/OS 427
 - CAM (Credit Application Manager) 608
 - caseIgnoreMatch 559, 561
 - caseIgnoreSubstringsMatch 559, 561
 - CCSID (coded character set identifier) 286
 - CCSID (Coded Character Set Identifier)
 - same as queue manager 15
 - CEDA 307
 - cepl 550
 - certificate revocation list (CRL) 1079
 - channel
 - auto-definition exit program 362
 - client connection 295
 - data-conversion exit 364
 - startup, data negotiation 342, 349
 - synchronizing 342
 - channel compression
 - using IBM MQ classes for Java 1068
 - using IBM MQ classes for JMS 992
 - channel definition 561
 - maximum message length 286
 - channel exit
 - compatible receive exit 359
 - reserving space in send exit 359
 - channel exits
 - assigning
 - with IBM MQ classes for Java 1064
 - with IBM MQ classes for JMS 979
 - auto-definition 362

- channel exits (*continued*)
 - message 360
 - message-retry 362
 - not written in Java
 - used with IBM MQ classes for Java 1066
 - used with IBM MQ classes for JMS 979
 - passing user data
 - with IBM MQ classes for Java 1066
 - with IBM MQ classes for JMS 981
 - receive 358
 - security 349
 - CLNTCONN and SVRCONN channels 357
 - send 358
 - specifying
 - classes for .NET 661
 - using 660
 - with IBM MQ classes for Java 1062
 - using a sequence
 - with IBM MQ classes for Java 1067
 - with IBM MQ classes for JMS 979, 981
 - with TLS 660
 - writing 660
 - written in Java
 - used with IBM MQ classes for Java 1063
 - used with IBM MQ classes for JMS 978
- channel-exit program 357
- channel-exit programs 341
 - channel definition structure, MQCD 344
 - data buffer 344
 - IBM MQ for HP-UX 363
 - IBM MQ for IBM i 347
 - IBM MQ for Windows 348
 - IBM MQ for z/OS 345
 - parameter structure, MQCXP 344
 - Windows client 348
 - writing and compiling 344
- ChannelAutoDefinition property 697
- ChannelAutoDefinitionEvent property 697
- ChannelAutoDefinitionExit property 698
- character strings 804, 806
- CharacterSet property 698, 735
- CICS
 - API-crossing exit 263
 - assembler language applications 455
 - calling the stub dynamically with IBM MQ for z/OS 427
 - COBOL applications 415
 - CSQCAPX 263
 - IBM MQ for z/OS support 259
 - on IBM i 406
 - preparing and running client applications 308
 - sample client programs 309
 - storage protection facility 264
 - with IBM MQ for AIX 386
 - with IBM MQ for HP-UX 397
- CICS (*continued*)
 - with IBM MQ for Solaris 412
 - with IBM MQ for z/OS 425
- CICS applications (non-z/OS)
 - CEDA 307
 - CICSENV.CMD file 307
 - environment specifics 307
 - MQSERVER 307
- CICS bridge, writing a message to 813
- CICS sample transaction 481
- CICS Transaction Server
 - running applications 1039
- CKTI transaction 248
- class path
 - settings 833
- classes 558
- classes for .NET 657
- classes, core
 - IBM MQ for Java 1087
 - restrictions and variations 1088, 1092
- classpath
 - settings 1029
- Clear method 786
- ClearErrorCodes method
 - MQDistributionList class 768
 - MQDistributionListItem class 776
 - MQGetMessageOptions class 763
 - MQMessage class 746
 - MQPutMessageOptions class 760
 - MQQueue class 726
 - MQQueueManager class 708
 - MQSession class 694
- ClearMessage method 746
- client
 - connection 851, 1022
- client (IBM MQ)
 - triggering support 234
- client applications
 - CICS and Tuxedo sample programs 309
 - preparing and running
 - CICS and Tuxedo 308
 - Microsoft Transaction Server 42, 310
 - MQ JMS 311
 - WebSphere Application Server 311
- client bindings connections 632, 1191
- client channel definition table
 - how it is used 297
 - using IBM MQ classes for Java 1047
 - using IBM MQ classes for JMS 982
 - where to find it 294
- client configuration file
 - .NET 650
- client connections
 - .NET managed 647
- client library file 292
- client mode connections 1021
- CLIENT_ACKNOWLEDGE 944
- client-connection channel
 - security exit 357
- clients
 - configuring queue manager 635, 852, 1040
 - programming
 - Java 1044
- clients (*continued*)
 - verifying 1040
- Close method 726, 768
- close, implicit operation 805
- CloseOptions property 698, 712, 764
- closing
 - resources using JMS 955
- cluster workload exit
 - sample 380
 - use of 379
 - writing and compiling 381
- clusters, use of
 - return codes 257
- CMQC.H header file 803
- CN 559, 561
- COBOL
 - CICS applications 415
 - on AIX 384
 - on HP-UX 395
 - on IBM i 405
 - on Linux 403
 - on Solaris 410
 - on Windows systems 415
 - support for 451
- code examples
 - classes for .NET 651
- code level tool 781
- coded character set identifier (CCSID) 286
- Coded Character Set Identifier (CCSID)
 - same as queue manager 15
- CodedCharSetId (CCSID) 15
- message data 16
- com.ibm.mq.headers.jar 1026
- com.ibm.mq.jar 1026
- com.ibm.mq.pcf.jar 1026
- COM+ 639
- CommandInputQueueName
 - property 698
- CommandLevel property 699
- commands using Windows 777
- commit
 - single-phase 222
 - two-phase 222
- Commit method 708
- committing changes 222
- common name 561
- compile cobol sample z/OS 576
- compiler commands 814
- compilers for IBM MQ platforms, overview 814
- compiling
 - for IBM MQ for AIX 382
 - for IBM MQ for HP-UX 393
 - C programs 393
 - COBOL programs 395
 - for IBM MQ for IBM i 404
 - for IBM MQ for Solaris
 - C programs 408
 - COBOL programs 410
 - for IBM MQ for Windows 413
 - for IBM MQ for z/OS 420
 - MQ for Linux
 - COBOL programs 403
- compiling exit on AIX 556
- compiling exit on Linux 556
- compiling exit on Solaris 556

- compiling exit on windows 556
 - compiling IBM MQ .NET programs 678
 - compiling sample programs
 - IBM i 816
 - z/OS 821
 - CompletionCode property
 - MQDistributionList class 765
 - MQDistributionListItem class 771
 - MQGetMessageOptions class 761
 - MQMessage class 731
 - MQPutMessageOptions class 758
 - MQQueue class 713
 - MQQueueManager class 699
 - MQSession class 691
 - components, installable services 326
 - compression of data 358
 - Configuration 645
 - configuration file
 - authorization service 330
 - IBM MQ classes for JMS 837
 - IBM MQ for Java 1033
 - configuration files
 - .NET 650
 - configuring
 - environment variables 833, 1029
 - Java 2 Security Manager 1037
 - Java security manager 845
 - JTA/JDBC coordination
 - other platforms 1074
 - Windows 1073
 - MQ resource adapter
 - inbound communication 1106
 - introduction 1103
 - outbound communication 1121
 - ResourceAdapter object 1104
 - queue manager for clients 852, 1040
 - your class path 833
 - your classpath 1029
 - configuring queue manager for clients 635
 - confirmation of arrival (COA) report 11
 - confirmation of delivery (COD) report 11
 - Connect method 708
 - connect, implicit operation 805
 - connecting to a client 295
 - connecting to a queue manager 92, 94, 654
 - connecting to a queue manager in Java 1047
 - connection
 - building 942
 - creating 942
 - modes 851
 - options 1022
 - queue managers 294
 - starting 942
 - connection endpoint lookup 550, 556
 - connection factories
 - creating
 - connection 942
 - connection factories
 - summary 930
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - destinations
 - summary 930
 - using a session 945
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - session 943
 - creating service components 329
 - connection factories (*continued*)
 - setting properties (*continued*)
 - using the IBM MQ JMS
 - extensions 932, 939
 - connection handle
 - returned from MQCONN 93
 - returned from MQCONNX 94
 - scope of 93
 - shared 97
 - thread independent 97
 - Connection interface 871
 - connection options 1021
 - connection pooling
 - example 1069
 - connection tag 267
 - connection type, defining 648, 1045
 - connectionFactory 1186
 - ConnectionFactory interface 871
 - ConnectionHandle property 699
 - ConnectionReference property 713, 765
 - connections
 - client bindings 632, 1191
 - managed client 632, 1191
 - server bindings 632, 1191
 - sharing
 - in JMS 990
 - ConnectionStatus property 699
 - ConnectOptions property 700
 - context
 - default 118
 - message 40
 - MQOPEN options 108
 - MQPUT options 118
 - Context field 114
 - conventions 805
 - Conversations, sharing 289
 - convert message data
 - MQGET 131, 164
 - converting the log file 863
 - copying messages 165
 - copyOut method 810
 - core classes
 - IBM MQ for Java 1087
 - restrictions and variations 1088, 1092
 - CorrelationId property 736, 771
 - CorrelationIdHex property 736, 772
 - CorrelId field 146
 - correlid, performance considerations 54
 - CorrelId, performance considerations 48
 - Count property 784
 - creating
 - connection 942
 - connection factories
 - summary 930
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - destinations
 - summary 930
 - using a session 945
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - session 943
 - creating service components 329
- CreationDateTime property 713
- Credit Application Manager (CAM) 608
- credit check sample (IBM MQ for z/OS) 604
- CSQ4BAA1 sample 590
- CSQ4BCA1 sample 590
- CSQ4BVA1 sample 590
- CSQ4CAC1 sample 596
- CSQ4CCB5 sample 613
- CSQ4CCC1 sample 596
- CSQ4CCG1 sample 592
- CSQ4CVB1 sample 608
- CSQ4CVB2 sample 608
- CSQ4CVB3 sample 612
- CSQ4CVB4 sample 612
- CSQ4CVB5 sample 613
- CSQ4CVC1 sample 596
- CSQ4CVD1 sample 601
- CSQ4CVD2 sample 602
- CSQ4CVD3 sample 602
- CSQ4CVD4 sample 603
- CSQ4CVD5 sample 604
- CSQ4TCD1 sample 601
- CSQ4TCD2 sample 602
- CSQ4TCD4 sample 603
- CSQ4TCD5 sample 604
- CSQ4TVD1 sample 601
- CSQ4TVD2 sample 602
- CSQ4TVD4 sample 603
- CSQ4TVD5 sample 604
- CSQCAPX API-crossing exit sample 265
- CSQCAPX sample 263
- CSQQTRMN transaction 248
- CurrentDepth property 714
- cursor, browse 106
- ## D
- data
 - application 9
 - compression 358
 - conversion 360
 - decompression 358
 - encryption 360
 - message 9
 - preparation 807
 - types 804
 - data conversion 294, 685
 - amqsvfc0 sample program 482
 - AMQSVFC4 sample program 482
 - amqsvfcx sample program 482
 - application 16
 - IMS bridge 69, 282
 - interface 365
 - message 164
 - MQGET 131, 164
 - MQXCNCV call 79
 - data conversion interface (DCI) 365
 - data definition files
 - copy files 71
 - header files 71
 - include files 71
 - macros 71
 - data in a message 9
 - data types
 - structures 79
 - data-conversion exit 364

- data-conversion exit (*continued*)
 - amqsvfc0 sample program 482
 - AMQSVFC4 sample program 482
 - amqsvfcx sample program 482
 - IMS bridge 69, 282
 - invoking 365
 - MQXCNCV call 79
 - writing
 - IBM i 367
 - UNIX systems 369
 - Windows NT 372
 - z/OS 368
 - datagram, putting and getting
 - C++ 799
 - DataLength parameter 133
 - DataLength property 732
 - DataOffset property 732
 - dataPointer method 811
 - date and time of messages 41
 - dead-letter (undelivered message) queue
 - handler 440
 - brief description 440
 - sample 489
 - sample to deal with messages on it 489
 - using 440
 - dead-letter queue
 - processing in Java 1055
 - dead-letter queue, writing a message to 812
 - DeadLetterQName field 241
 - DeadLetterQueueName property 700
 - debugging programs 433
 - declaring parameters 805
 - decompression of data 358
 - default connection pool 1069
 - multiple components 1071
 - default context 118
 - DefaultInputOpenOption property 714
 - DefaultPersistence property 714
 - DefaultPriority property 714
 - DefaultTransmissionQueueName property 700
 - DEFBIND attribute
 - inquiring on a queue 256
 - opening a queue 254
 - defining alias for queue manager 102
 - defining channels
 - using MQSC 296
 - defining connection type 648, 1045
 - DefinitionType property 715
 - DefInputOpenOption attribute 107
 - DefPriority attribute 17
 - DepthHighEvent property 715
 - DepthHighLimit property 715
 - DepthLowEvent property 715
 - DepthLowLimit property 716
 - DepthMaximumEvent property 716
 - Description property 700, 716
 - design and programming using IBM MQ
 - automation classes for ActiveX 682
 - designing applications that access non-ActiveX applications 682
 - Destination interface 871
 - destinations
 - creating
 - summary 930
 - destinations (*continued*)
 - creating (*continued*)
 - using a session 945
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - setting properties
 - summary 930
 - using a URI 945
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - differences due to environment 1092
 - directories, installation 1028
 - Disconnect method 708
 - disconnect, implicit operation 805
 - disconnecting from a queue manager 99, 654
 - disconnecting from a queue manager in Java 1047
 - disposition options, message 1005
 - distributed transaction 636
 - distribution list, putting messages to 802
 - distribution lists 122
 - opening 123
 - platform dependency 1091
 - putting messages to 126
 - DistributionList property 772
 - DistributionLists property 701
 - DPUT sample program 802
 - DTC 639
 - DUPS_OK_ACKNOWLEDGE 944
 - durable topic subscribers 950
 - dynamic binding 322, 328
 - dynamic linking of MQI calls for IBM MQ for z/OS 427
 - dynamic queue
 - closing temporary queue 110
 - creating 109
 - dynamic queues, shared queues 268
 - dynamic XA resource management
 - structure 232
 - DynamicQName field 109
 - DynamicQueueName property 716
- E**
- EBCDIC newline character
 - conversion 165
 - ECB (event control block) 160
 - Encoding field 16
 - Encoding property 737
 - encryption
 - in send exit 359
 - encryption of messages 341
 - endmqdmn 675
 - environment dependencies
 - functions not with all platforms 1090
 - distribution lists 1091
 - MQGetMessageOptions
 - fields 1090
 - MQMD fields 1091
 - MQPutMessageOptions
 - fields 1091
 - environment dependencies (*continued*)
 - functions not with all platforms (*continued*)
 - MQQueueManager begin()
 - method 1090
 - MQQueueManager
 - constructor 1090
 - IBM MQ for Java 1087
 - restrictions and variations 1088
 - MQGMO_* values 1088
 - MQPMO_* values 1089
 - MQPMRF_* values 1089
 - MQRO_* values 1089
 - z/OS 1089
 - environment differences 1092
 - environment variable
 - MQ_CONNECT_TYPE 99
 - environment variables 295, 833, 1029
 - MQCHLLIB 294
 - MQCHLTAB 294
 - MQSERVER 294
 - environments for IBM MQ for z/OS 258
 - error
 - handling 1060
 - logging 863
 - IBM MQ classes for JMS 957
 - error handling 687
 - Error handling 782
 - error handling in the samples 792
 - error messages 635
 - errors
 - classes for .NET 656
 - dead-letter (undelivered message) queue 440
 - dealing with failure of a call 436
 - report message 438
 - undelivered message queue 440
 - errors on parameter passing 689
 - escaping, in URI 1196
 - event control block 160
 - example code
 - classes for .NET 651
 - example output file from code level tool 781
 - examples
 - custom encapsulated message-writing code 810
 - declaration and use conventions 805
 - headers 805
 - ImqDeadLetterHeader class 810
 - invoking the cluster workload
 - exit 380
 - manipulating binary strings 804
 - preparing message data 807
 - retrieving items within a message 808
 - retrieving messages into a fixed area of storage 811
 - sample programs 798
 - DPUT (imqdput.cpp) 802
 - SGET (imqsgget.cpp) 802
 - writing a message to the CICS bridge 813
 - writing a message to the dead-letter queue 812
 - writing a message to the IMS bridge 812

- examples (*continued*)
 - writing a message to the work header 814
- exception listeners 957
- exception report options, message 1005
- exceptions
 - IBM MQ classes for JMS 957
- exclusive access to a queue 107
- execute method 786
- Execute method
 - Command parameter 787
 - OptionsBag parameter 787
 - QueueManager parameter 787
 - ReplyBag parameter 787
 - ReplyQ parameter 787
 - RequestQ parameter 787
- execution key of CICS programs 264
- exit module 556
- exit programs 263
- exit, cluster workload
 - sample 380
 - use of 379
 - writing and compiling 381
- Expiry property 738
- expiry report 11
- explicit programming model 638
- external syncpoint
 - interfaces 231
 - restrictions 232
 - X/Open XA interface 231

F

- failure of the IBM MQ automation classes
 - for ActiveX script 778
- FASTPATH binding 94
 - environment variable 99
- feedback codes, IMS bridge 65, 278
- Feedback field 13
- Feedback property 738, 772
- fields
 - AlternateUserId 108
 - AppIdentityData 40
 - AppOriginData 40
 - BackoutCount 39
 - Context 114
 - CorrelId 146
 - DeadLetterQName 241
 - DynamicQName 109
 - Encoding 16
 - Feedback 13
 - Format 15
 - GroupId
 - match options 146
 - MQMO 146
 - InitiationQName 239
 - MsgId 146
 - Persistence 38
 - Priority 17
 - PutAppName 40
 - PutAppType 40
 - PutDate 40
 - PutMsgRecFields 115
 - PutMsgRecOffset 115
 - PutMsgRecPtr 115
 - PutTime 40
 - RecsPresent 114

- fields (*continued*)
 - ReplyToQ 39
 - ReplyToQMGr 39
 - ResolvedQMGrName 114
 - ResolvedQName 114
 - ResponseRecOffset 115
 - ResponseRecPtr 115
 - StrucId 114
 - UserIdentifier 40
 - Version 114
 - WaitInterval 131, 159
- FirstDistributionListItem property 765
- format
 - control information 15
 - message data 15
- Format field 15
- Format property 738
- formatLog utility 863
- Frommessage method 787
- FromMessage method
 - Message parameter 787

G

- get (sample for IBM MQ for z/OS) 587
- get method 811
- Get method 726
- getting
 - a particular message 146
 - message when the length is unknown 168
 - messages 129
 - options 128
- getting a datagram, sample program
 - C++ 799
- getting started
 - IBM MQ classes for JMS 827
 - IBM MQ for Java 1022
- getting trigger messages
 - trigger queue 247
- Global Assembly Cache 639
- global transaction 636
- GMT (Greenwich Mean Time) 41
- group attach 4
- group batch attach 4
- GroupId property 739, 773
- GroupIdHex property 739, 773
- GroupStatus field
 - MQGMO structure 132

H

- handle
 - scope of connection handle 93
 - using object handle 100
- handling
 - errors 1060
 - classes for .NET 656
 - headers, in Java 1053
 - messages 655, 1051
- HardenGetBackout attribute 39, 438
- HardenGetBackout property 717
- Hconn
 - shared 97
 - thread independent 97

- header
 - handling, Java 1053
- header example 805
- header files
 - CMQC.H 803
 - IMQI.HPP 803
 - IMQTYPE.H 803
- headers
 - creating classes 1058
 - finding 1056
 - printing in Java 1054, 1056
 - skipping over in Java 1054
- headers, message 875, 1172
- HELLO WORLD sample program
 - C++ 799
- how triggering works 234
- HP-UX
 - Address space models 398

I

- IBM i compilers
 - IBM ILE C++ 816
 - VisualAge C++ 816
- IBM i compiling 816
- IBM ILE C++ 816
- IBM JMS extensions
 - introduction
 - general 824
 - using 932
- IBM MQ classes for JMS
 - benefits 824
 - configuration file 837
 - introduction
 - general 824
- IBM MQ Client for HP Integrity NonStop Server
 - notification of message arrival 160
 - using signaling 160
- IBM MQ for AIX
 - amqisetx 518
 - amqmssetx 518
 - amqsseta 518
 - build TUXEDO server
 - environment 525, 526
 - building your application 382
 - CICS support 386
 - sample programs 458
 - set sample 518
 - syncpoints 228
 - triggering using samples 514
 - TUXEDO samples 524
- IBM MQ for HP-UX
 - amqisetx 518
 - amqmssetx 518
 - amqsseta 518
 - build TUXEDO server
 - environment 527, 528
 - building your application 393
 - channel-exit programs 363
 - CICS support 397
 - sample programs 458
 - set sample 518
 - triggering using samples 514
 - TUXEDO samples 524
- IBM MQ for IBM i
 - AMQZSTUB 404

- IBM MQ for IBM i *(continued)*
 - building your application 404
 - channel-exit programs 347
 - compiling 404
 - CRTCMOD 404
 - disconnecting from queue manager 99
 - linking 404
 - sample program
 - using triggering 516
 - SQL programming
 - considerations 407
 - syncpoint considerations with CICS for IBM i 227
 - syncpoints 228, 233
 - trigger monitors 251
 - triggering Java applications 251
- IBM MQ for Java
 - configuration file 1033
 - connection options 1021
 - introduction 1021
- IBM MQ for Solaris
 - build TUXEDO server
 - environment 529, 530
 - building your application 407
 - C compiler 408
 - CICS support 412
 - sample programs 458
- IBM MQ for Windows
 - amqrs핀.dll 539
 - amqssp인.c 539
 - authentication
 - Kerberos 419
 - NTLM 418
 - build TUXEDO server
 - environment 531, 534
 - building your application 413
 - channel-exit programs 348
 - context acceptor, security exit 418
 - context initiator, security exit 418
 - Kerberos
 - authentication 419
 - NTLM authentication 418
 - object code, security exit 539
 - principal, security exit 419
 - sample programs 458
 - security exit
 - object code 539
 - source code 539
 - source code, security exit 539
 - syncpoints 228
 - triggering using samples 514
 - TUXEDO sample makefile 532, 535
 - TUXEDO samples 524
 - ubbstxcn.cfg example 531, 534
 - IBM MQ for z/OS
 - building your application 420
 - channel-exit programs 345
 - CMQA 74
 - CMQDLHA 74
 - CMQDXPA 74
 - CMQGMOA 74
 - CMQIIHA 74
 - CMQMDA 74
 - CMQODA 74
 - CMQPMOA 74
 - CMQTMA 74
 - IBM MQ for z/OS *(continued)*
 - CMQTMCA 74
 - CMQXA 74
 - CMQXPA 74
 - CMQXQHA 74
 - CSQQSTUB 426
 - requesting no backout of MQGET 161
 - using signaling 160
 - IBM MQ resource adapter
 - limitations 1096
 - IBM Tivoli Directory Server 6.3 556
 - ibm-amqChannelName 559, 561
 - ibm-amqClientChannelWeight 559
 - ibm-amqConnectionAffinity 559
 - ibm-amqConnectionName 559, 562
 - Ibm-amqDescription 559, 561
 - ibm-amqHeaderCompression 559, 566
 - ibm-amqHeartBeatInterval 559, 564
 - ibm-amqIsClientDefault 565
 - ibm-amqKeepAliveInterval 559, 565
 - ibm-amqLocalAddress 559, 562
 - ibm-amqMaximumMessageLength 559, 565
 - ibm-amqMessageCompression 559, 566
 - ibm-amqModeName 559, 562
 - ibm-amqPassword 562
 - ibm-amqQueueManagerName 559
 - ibm-amqSendExitUserData 559
 - ibm-amqSendExitUserName 559
 - ibm-amqSslCipherSpec 559
 - ibm-amqSslPeerName 559
 - ibm-
 - amqTransactionProgramName 559
 - ibm-amqUserID 559
 - ibm-wmqSecurityExitName 559
 - ibm-amqQueueManagerName 562, 565
 - ibm-amqReceiveExitName 568
 - ibm-amqReceiveExitUserData 559, 567
 - ibm-amqSecurityExitName 563
 - ibm-amqSecurityExitUserData 563
 - ibm-amqSendExitName 567
 - ibm-amqSendExitUserData 567
 - ibm-amqSharingConversations 559, 565
 - ibm-amqSslCipherSpec 563
 - ibm-amqSslPeerName 563
 - ibm-amqTransactionProgramName 564
 - ibm-amqTransportType 559, 565
 - ibm-amqUserID 564
 - ibm-wmqClientChannelWeight 564
 - ibm-wmqConnectionAffinity 564
 - implicit operations 805
 - implicit programming model 638
 - IMQI.HPP header file 803
 - IMQObjectTrigger 675
 - IMQTYPE.H header file 803
 - IMS
 - adapter 262
 - building application for IBM MQ for z/OS 426
 - calling the stub dynamically with IBM MQ for z/OS 427
 - closing objects 100
 - mapping IBM MQ messages to transactions 64, 277
 - MQITS_ARCHITECTED
 - constant 284
 - IMS *(continued)*
 - support for 259
 - using MQI calls 59, 272
 - using OTMA 284
 - using syncpoints 59, 271
 - writing IBM MQ applications 58, 270
 - IMS bridge
 - alternate PCB 68, 281
 - data conversion 69, 282
 - feedback codes 65, 278
 - IMS commands 63, 275
 - LLZZ data segment 69, 281
 - mapping IBM MQ messages to transactions 64, 277
 - message segmentation 69, 281
 - NAK 63, 275
 - reply messages 68, 280
 - sense codes 65, 278
 - undelivered messages 63, 275
 - writing applications 63, 275
 - IMS bridge, writing a message to 812
 - IMS commands, IMS bridge 63, 276
 - index queues
 - shared queues 268
 - index, queue 54
 - InhibitEvent property 701
 - InhibitGet property 717
 - InhibitPut property 717
 - initial state for objects 804
 - initialContextFactory 1186
 - initialization 317
 - initiation queue
 - example to create one 239
 - what it is 235
 - initiation queue, shared 268
 - InitiationQName field 239
 - InitiationQueueName property 718
 - inquire and set 1061
 - classes for .NET 657
 - inquiring about attributes
 - IBM MQ for IBM i sample program 496
 - IBM MQ for UNIX sample program 496, 497
 - IBM MQ for Windows sample program 496
 - IBM MQ for z/OS sample 596
 - using MQINQ 218
 - inquiring about properties of a message handle
 - IBM MQ for IBM i sample program 497
 - IBM MQ for Windows sample program 497
 - installable service
 - authorization service 330
 - component data 315
 - component entry-points 315
 - components 315, 327
 - configuring services 322, 328
 - functions 315, 327
 - initialization 317, 328
 - multiple components 321
 - name service 318
 - name service interface 320
 - return information 315

- installation
 - classes for .NET 632
 - directories 1028
- Installation and Configuration 553
- installation verification test
 - IBM MQ classes for JMS
 - point-to-point 853
 - IBM MQ classes for JMS
 - publish/subscribe 857
- installation verification test (IVT)
 - program
 - IBM MQ resource adapter 1135
- installing
 - IBM MQ classes for JMS 830
 - IBM MQ for Java 1026
 - MQ resource adapter 1099
- interface, programming
 - .NET 630
- interfaces
 - JMS 871
- interfaces to external syncpoint
 - managers 231
- introduction
 - for programmers
 - using Java 1043
 - IBM MQ classes for JMS 827
 - IBM MQ for Java 1021, 1022
- Introduction 550
- invoking data-conversion exit 365
- IsConnected property 701
- IsOpen property 702, 718, 766
- item
 - description 808
 - retrieving from a message 808
- item property 783
- ItemIndex parameter
 - Item property 784
- itemtype method 788
- ItemType method
 - ItemIndex parameter 788
 - ItemType parameter 788
 - Selector parameter 788
- IVT (installation verification test
 - program)
 - IBM MQ resource adapter 1135
- IVT (installation verification test)
 - IBM MQ classes for JMS
 - point-to-point 853
 - IBM MQ classes for JMS
 - publish/subscribe 857
- IVTRun script 853, 861
- IVTSetup script 853, 861
- IVTTidy script 853, 861

J

- J2EE Connector Architecture (JCA)
 - IBM MQ resource adapter 1092
- jar files 830, 1026
- Java
 - introduction for programmers 1043
- Java 2 Security Manager, running
 - applications under 1037
- Java EE 824
- Java programs
 - triggering on IBM MQ for IBM i 251

- Java security manager, running
 - applications under 845
- JCA (J2EE Connector Architecture)
 - IBM MQ resource adapter 1092
- JCL (Job Control Language)
 - batch 421
 - CICS and IBM MQ for z/OS 425
 - IMS 426
 - XPLINK 423
- JDBC coordination 1073
- JMS 824
 - administered objects 871
 - applications, writing 871
 - benefits 824
 - interfaces 871
 - mapping of fields at send or
 - publish 887
 - mapping with MQMD 884
 - messages 875
 - persistent 970
 - model 871
 - persistent messages 970
- JMS_IBM_RETAIN 995
- JMSAdmin utility 861
- JMSCorrelationID header field 875
- JNDI
 - retrieving administered objects 931
- journaling 360
- JSSE
 - for TLS support 971, 1077
 - making changes to the keystore or
 - truststore
 - using IBM MQ classes for
 - Java 1081
 - using IBM MQ classes for
 - JMS 975
- JTA/JDBC coordination
 - configuring
 - other platforms 1074
 - Windows 1073
 - introduction 1073
 - known problems 1076
 - limitations 1076
 - using 1075

K

- keystore, making changes
 - using IBM MQ classes for Java 1081
 - using IBM MQ classes for JMS 975

L

- language considerations
 - attributes 803
 - binary strings 804
 - character strings 804
 - data types 804
 - header files 803
 - methods 803
 - notational conventions 805
 - using C from C++ 805
- languages supported
 - clients 291
- large messages
 - reference messages 152

- large messages (*continued*)
 - segmented messages 152
- LDAP 561
- LDAP attributes 559
- LDAP schemas 558
- libraries
 - IBM MQ classes for JMS 835
 - Java Native Interface (JNI) 835
- libraries to use
 - with IBM MQ for AIX 383
 - with IBM MQ for Linux 399
 - with IBM MQ for Windows 414
- libraries, IBM MQ classes for Java 1030
- library file, client 292
- library files 80
- lifelines, publish/subscribe 201
- link libraries 814
- linked exceptions 957
- linking 814
 - for IBM MQ for AIX 382
 - for IBM MQ for HP-UX 393
 - for IBM MQ for IBM i 404
 - for IBM MQ for Windows 413
 - for IBM MQ for z/OS 420
- linking with IBM MQ MQI client code
 - C applications 292
 - C++ applications 293
 - Visual Basic applications 293
- LLZZ data segment, IMS bridge 69, 281
- local administration
 - support for application programs 5
- local publications
 - suppressing 950
- local transactions
 - JMS 943
- LOCALADDRESS
 - specifying a range of ports 991
- LocalEvent property 702
- log file
 - converting 863
- logging errors 863
 - IBM MQ classes for JMS 957
- looking at a message 165

M

- mail manager sample application (IBM
 - MQ for z/OS) 597
- managed client connections 632, 1191
- managed mode 636, 638
- manipulating strings
 - example 804
 - introduction 806
- map message 875, 1172
- marked messages 171
- marking messages as browsed 171
- MatchOptions field
 - MQGMO structure 132
- MatchOptions property 762
- MaximumDepth property 718
- MaximumHandles property 702
- MaximumMessageLength property 702,
 - 718
- MaximumPriority property 703
- MaximumUncommittedMessages
 - property 703
- MaxMsgLength attribute 116, 133

- MaxPriority attribute 17
- message
 - backed out 39
 - body 875, 1172
 - body types 949
 - browsing 165
 - browsing and removing 168
 - browsing in logical order 168
 - browsing when message length unknown 168
 - confirm arrival 11
 - confirm delivery 11
 - context
 - MQOPEN options 108
 - MQPUT options 118
 - types 40
 - copying 165
 - creating 9
 - data 9
 - data conversion
 - considerations 16
 - MQGET 164
 - data format 15
 - delivery, asynchronous 950
 - design 46
 - encryption 341
 - error 635
 - exception 11
 - expiry 11
 - getting 129
 - getting a particular 146
 - greater than 4 MB 152
 - groups 36
 - handling 655, 1051
 - headers 875, 1172
 - logical ordering 134
 - looking at 165
 - maximum length 286
 - message body 893
 - negative action notification 11
 - notification of arrival 160
 - order of retrieval from a queue 133
 - persistence 38
 - physical ordering 134
 - positive action notification 11
 - priority 17, 133
 - properties 875, 1172
 - putting 111
 - putting one 121
 - reference 157
 - removing after browsing 168
 - reply, IMS bridge 68, 280
 - report 438
 - retry sending 439
 - sample to deal with those on
 - dead-letter queue 489
 - segmentation 153
 - segmented 37
 - selectors 875, 950
 - selectors and SQL 27
 - signaling 160
 - structure 9
 - trigger 235, 252
 - types 10, 875, 1172
 - undeliverable, IMS bridge 65, 277
 - undelivered 439
 - undelivered, sample to handle 489
- message (*continued*)
 - use of types 10
 - variable length, performance considerations 54
 - waiting for 159
- message buffers
 - application (manual) 807
 - system (automatic) 807
- message channel agent
 - initiation 349, 358
 - termination 349, 358
- Message Consumer
 - poison messages 956
- message content
 - printing in Java 1056
- message data conversion, MQGET 131, 164
- message data preparation 807
- message exit program 360
 - overview 342
- message handler sample (IBM MQ for z/OS) 617
- message headers
 - CICS bridge header 813
 - dead-letter header 812
 - IMS bridge header 812
 - work header 814
- Message interface 871
- message items
 - description 808
- message object properties 964
- message order 208
- message persistence, performance considerations 54
- message properties 18
 - classes for .NET 656
 - setting 118
- Message Queue Interface
 - calls 76
 - dealing with failure of a call 436
 - library files 80
 - structures 79
 - stub programs 80
 - using System/390 assembler 455
- Message Queue Interface (MQI) 286
- message segmentation, IMS bridge 69, 281
- message-retry exit program 362
- MessageConsumer interface 871
- MessageConsumer object 950
- MessageData property 740
- MessageDeliverySequence property 719
- MessageFlags property 740
- MessageId property 740, 773
- MessageIdHex property 741, 774
- MessageLength property 733
- MessageProducer interface 871
- MessageProducer object 949
- messages
 - AMQP 1172
 - JMS 875
 - mapping between JMS and IBM MQ 879
 - mapping to IMS transaction types 64, 277
 - persistent messages, effect on performance 48
- messages (*continued*)
 - persistent, JMS 970
 - placing on named queue, example C++ 802
 - poison
 - Application Server Facilities 1006
 - JMS 956
 - putting to a distribution list, example 802
 - reading 808
 - receiving 950
 - retrieving from named queue, example C++ 802
 - retrieving in correct order 267
 - selecting 875, 950
 - sending 949
 - variable length 48
 - writing
 - to the CICS bridge 813
 - to the dead-letter queue 812
 - to the IMS bridge 812
 - to the work header 814
- MessageSequenceNumber property 741
- MessageType property 742
- messaging domains 871
- method signatures 803
- methods 807
 - detailed description
 - Add method 785
 - Addinquiry method 786
 - Clear method 786
 - execute method 786
 - Frommessage method 787
 - itemtype method 788
 - Remove method 788
 - selector method 789
 - Tomessage method 790
 - Truncate method 790
- Microsoft Transaction Server (MTS)
 - preparing and running client applications 42, 310
- migrating applications to use shared queues 269
- model queue 109
- modes
 - connection 851
- Monitor program 675
- MQ
 - messages 879
 - name service interface (NSI) 318
 - Object Model 796
 - security enabling interface (SEI) 330
- MQ applications
 - planning 43
- MQ Automation Classes for Activex 681
- MQ automation classes for ActiveX
 - failure 778
- MQ automation classes for ActiveX interface 688
- MQ automation classes for ActiveX Reference 688
- MQ data conversion interface 365
- MQ for Linux
 - building your application 398
 - 31-bit 400
 - 32-bit 401

- MQ for Linux *(continued)*
 - building your application *(continued)*
 - 64-bit 402
- MQ JMS
 - preparing and running client applications 311
- MQ JMS extensions
 - introduction
 - general 824
 - using 939
- MQ Message Descriptor (MQMD) 879
 - mapping with JMS 884
- MQ object
 - closing 110
 - opening 100
 - process definition
 - create 240
- MQ on UNIX
 - sample programs 458
 - syncpoints 228
- MQ on UNIX systems 514
 - TUXEDO 524
- MQ resource adapter
 - configuration
 - inbound communication 1106
 - introduction 1103
 - outbound communication 1121
 - ResourceAdapter object 1104
 - Inbound poison message
 - handling 1106
 - installation 1099
 - installation verification test (IVT)
 - program 1135
 - introduction 1092
 - other required documentation 1103
 - properties
 - ActivationSpec object 1106
 - ConnectionFactory object 1121
 - Queue or Topic object 1121
 - WebSphere Application Server, using with 1097
- MQ_CONNECT_TYPE 99
- MQ_MSG_HEADER_LENGTH 117
- MQ_PRECONNECT_EXIT 374
- MQ*_DEFAULT values
 - with IBM MQ for AIX 449
- MQBACK 287
- MQBag class 782
- MQCA_* values 219
- MQCD, channel definition structure 344
- MQCHLLIB
 - how it is used 294
- MQCHLTAB
 - how it is used 294
- MQCLOSE
 - call parameters 110
 - closing a queue 110
- MQCMIT 287
- MQCNO_FASTPATH_BINDING
 - variations by environment 1089
- MQCONN 294
 - call parameters 93
- MQCONNX 94, 294
 - environment variable 99
 - options
 - FASTPATH binding 94
 - shared connection 97
- MQCONNX *(continued)*
 - options *(continued)*
 - STANDARD binding 94
 - thread independent connection 97
- MQCONNX call
 - Using 291
- MQCONNXAny call
 - use in Visual Basic 450
- MQCXP structure
 - ExitReason field 359
 - ExitSpace field 359
- MQCXP, channel exit parameter
 - structure 344
- MQDH 117
- MQDISC
 - when to use 99
- MQDistributionList class 763
- MQDistributionListItem class 769
- MQDLH 117, 440
 - example 1055
- MQEnvironment 648, 1045
- MQGET
 - backing out changes 222
 - call parameters 129
 - committing changes 222
 - data conversion 164
 - increase speed of 152
 - message data conversion 131
 - order of message retrieval 133
 - to get a specific message 146
 - unknown message length 168
 - waiting for messages 159
 - when it fails 172
 - when to use 128
- MQGetMessageOptions class 761
- MQGetMessageOptions fields
 - platform dependency 1090
- MQGMO_*
 - ACCEPT_TRUNCATED_MSG 132
- MQGMO_* values
 - variations by environment 1088
- MQGMO_BROWSE_*
 - MSG_UNDER_CURSOR 168
- MQGMO_CONVERT 164
- MQGMO_MARK_SKIP_BACKOUT 438
 - explanation 161
- MQGMO_MSG_UNDER_CURSOR 168
- MQGMO_WAIT 159
- MQHeader 1053
 - example 1057
 - implementing 1058
- MQHeaderIterator 1053
 - example 1054, 1057
- MQHeaderList 1053
 - example 1056
- MQI
 - application in client environment 294
- MQI (Message Queue Interface) 286
 - calls 76
 - client library files 80
 - dealing with failure of a call 436
 - IMS applications 59, 272
 - library files 80
 - structures 79
 - stub programs 80
 - using System/390 assembler 455
- MQI (message-queuing interface)
 - local administration support 5
- MQI calls
 - MQINQ 256
 - MQOPEN 254
 - MQPUT 256
 - MQPUT1 256
 - MQSET 257
- MQIA_* values 219
- MQIIH 63, 275
- MQINQ 287
 - call parameters 220
 - use of selectors 219
 - when it fails 221
- MQINQ call 256
- MQIVP
 - listing 1041
 - sample application 1040
 - tracing 1041
- MQMD (MQSeries Message Descriptor) 879
- MQMD fields
 - platform dependency 1091
- MQMessage 655, 1051
- MQMessage class 728
- MQMF_SEGMENTATION_ALLOWED
 - option 254
- MQMT_* values 10
- MQOD 102
- MQOO_* values 105
 - MQOO_BIND_AS_Q_DEF option 254
 - MQOO_BIND_NOT_FIXED option 254
 - MQOO_BIND_ON_OPEN option 254
- MQOPEN
 - call parameters 101
 - MQOO_* values 105
 - object handle 100
 - resolving local queue names 109
 - using MQOD 102
 - using options parameter 105
- MQOPEN call
 - description 254
- MQPMO_* values
 - variations by environment 1089
- MQPMRF_* values
 - variations by environment 1089
- MQPUT
 - backing out changes 222
 - call parameters 112
 - committing changes 222
 - context information 118
 - if it fails 127
 - quiescing queue manager 114
 - sync pointing 114
- MQPUT and MQPUT1, performance considerations 48, 55
- MQPUT call 256
- MQPUT1
 - call parameters 121
 - if it fails 127
 - performance 111
- MQPUT1 call 256
- MQPutMessageOptions class 758
- MQPutMessageOptions fields
 - platform dependency 1091
- MQQueue 655, 1051
- MQQueue class 709

MQQueueManager 1050
 MQQueueManager begin() method
 platform dependency 1090
 MQQueueManager class 694
 MQQueueManager constructor
 platform dependency 1090
 MQRC_*
 SECOND_MARK_
 NOT_ALLOWED 162
 MQRC_* return codes
 CLUSTER_EXIT_ERROR 257, 381
 CLUSTER_EXIT_LOAD_ERROR 257
 CLUSTER_PUT_INHIBITED 257
 CLUSTER_RESOLUTION
 _ERROR 257
 CLUSTER_RESOURCE_ERROR 257
 NO_DESTINATIONS
 _AVAILABLE 257
 STOPPED_BY_CLUSTER_EXIT 257
 MQReceiveExit 660
 MQRFH2
 analyzing, example 1057
 MQRFH2 header 880
 MQRMH, reference-message header 360
 MQRMIXASwitch 232
 MQRMIXASwitchDynamic 232
 MQRO_* values
 variations by environment 1089
 mqs.ini file
 role in connecting a client 294
 MQSecurityExit 660
 MQSendExit 660
 MQSeries for OS/2 Warp
 sample programs 458
 MQSERVER
 how it is used 294
 relationship with client channel
 definition table 297
 MQSession Class 690
 MQSET
 attribute list 221
 call parameters 221
 use of selectors 219
 MQSET call 257
 MQTM (trigger message) 250
 MQTMC (trigger message,
 character) 250
 MQTMC2 (trigger message,
 character) 250
 MQTopic 1051
 MQWXP data structure 379
 MQXCNCV data-conversion call 79
 MQXCP_VERSION_5, of MQCXP
 structure 359
 MQXQH 117
 MQXQH, transmission header 360, 362
 MQXR_INIT, ExitReason value 359
 MQXR_XMIT, ExitReason value 359
 MS DTC 636
 MsgDeliverySequence attribute 134
 MsgHandle field
 MQGMO structure 132
 MsgId field 146
 msgid, performance considerations 54
 MsgId, performance considerations when
 using 48

MsgToken field
 MQGMO structure 132
 multiple queue managers 297
 multiple service components 321
 multithreaded programs 658
 multithreaded programs in Java 1062

N

Name property 703, 719
 name resolution 102
 name service
 interface (NSI) 318
 namelist
 opening 100
 sample application 612
 negative action notification (NAN)
 report 11
 NET managed clients
 programming 647
 NextDistributionListItem property 774
 NmqiXaResource 638
 nondurable topic subscribers 950
 Nonpersistent messages, tuning in
 AIX 151
 notational conventions, example 805
 notification of message arrival 160
 NSI (IBM MQ name service
 interface) 318

O

OAM
 refreshing after changing a user's
 authorization 326
 object
 closing 110
 descriptor 102
 opening 100
 process definition 240
 using handle 100
 object access methods 689
 object attributes 803
 object authority manager 318
 ObjectHandle property 719
 objects
 administered 871
 retrieving from JNDI 931
 message 875, 1172
 objects, initial state 804
 Offset property 742
 OID 558
 Open method 768
 open, implicit operation 805
 openFor method 805
 opening an IBM MQ object 100
 OpenInputCount property 719
 OpenOptions property 720, 766
 OpenOutputCount property 720
 OpenStatus property 720
 operations on queue managers
 Java 1045
 options
 connection 1022
 Options field
 MQGMO structure 131

Options field (*continued*)
 MQPMO structure 114
 Options parameter
 Options Property 785
 Options parameter (MQOPEN call) 105
 options property 785
 Options property 759, 762
 order of message retrieval 133
 OriginalLength property 742
 OSGi support 862, 1031
 OTMA sense codes 65, 278
 Overview of exit and schema 554

P

parameter passing 689
 parameters
 DataLength 133
 declaring 805
 Options 105
 passing 803
 security exit 355
 passing parameters 803
 pasteln method 810
 PCF
 handling in Java 1059
 pCompCode 374
 performance
 advantages of using MQPUT1 48
 application design, impact on 48
 application programs
 message length, effects on
 performance 48
 CorrelId, effect on 48
 message persistence, effect on 48
 messages
 message length, effects on
 performance 48
 MQGET and buffer size 133
 MQGET for a particular message 147
 MQPUT1 111
 MsgId, effect on 48
 nonpersistent messages in AIX 151
 performance
 message length, effects on 48
 persistent messages 38
 sync points, effects on 48
 threads, effect on 48
 performance considerations 55, 57
 PerformanceEvent property 704
 Persistence field 38
 Persistence property 743
 persistent messages, effect on
 performance 48
 persistent messages, JMS 970
 pExitParms 374
 PL/I
 support for 456
 placing messages on named queue,
 example
 C++ 802
 planning an IBM MQ application 43
 platform differences 1092
 Platform property 704
 point-to-point installation verification
 test 853
 point-to-point messaging 871

- poison messages
 - Application Server Facilities 1006
 - JMS 956
 - ports, specifying a range for client connections
 - IBM MQ classes for JMS 991
 - IBM MQ for Java 1049
 - positive action notification (PAN) report 11
 - ppConnectOpts 374
 - pQMGrName 374
 - pReason 374
 - preconnect stanza 376
 - Preconnection exit 374
 - preparing client applications
 - CICS 308
 - Microsoft Transaction Server 42, 310
 - MQ JMS 311
 - Tuxedo 308
 - WebSphere Application Server 311
 - preparing message data
 - example 807
 - introduction 807
 - preparing to run the samples 792
 - prerequisite software 827
 - classes for .NET 632
 - IBM MQ for Java 1022
 - PreviousDistributionListItem property 774
 - primary initialization 317
 - primary termination 317
 - print message (sample for IBM MQ for z/OS) 592
 - Priority field 17
 - priority in messages 17
 - Priority property 743
 - problem delivering a message, overview 39
 - problem determination
 - application design considerations 48
 - performance 51
 - searching for messages, performance effects 48
 - problem determination, use of report message 438
 - problems, solving 635
 - IBM MQ classes for JMS
 - general 863
 - point-to-point installation verification test 853
 - publish/subscribe installation verification test 857
 - IBM MQ for Java 1041
 - process definition 306
 - process definition object
 - example to create one 239
 - opening 100
 - triggering prerequisite 240
 - what it is 235
 - processes
 - accessing
 - in Java 1050
 - ProcessName 250
 - ProcessName attribute 240
 - ProcessName property 721
 - programming
 - bindings connection
 - Java 1044
 - client connections
 - Java 1044
 - compiling
 - .NET 678
 - connections 647, 1043
 - managed client connections 647
 - multithreaded 658
 - Java 1062
 - tracing
 - .NET 679
 - writing 647
 - writing Java programs 1043
 - z/OS 821
 - Programming hints and tips 683
 - programming interface
 - .NET 630
 - programs
 - JMS, writing 871
 - running 1086
 - properties
 - connection factories
 - setting by using the IBM JMS extensions 932
 - setting by using the IBM MQ JMS extensions 939
 - setting summary 930
 - destinations
 - setting by using a URI 945
 - setting by using the IBM JMS extensions 932
 - setting by using the IBM MQ JMS extensions 939
 - setting summary 930
 - detailed description
 - Count property 784
 - item property 783
 - options property 785
 - message 875, 1172
 - JMS_IBM_MQMD_* 964
 - message properties 18
 - MQ resource adapter
 - ActivationSpec object 1106
 - ConnectionFactory object 1121
 - Queue or Topic object 1121
 - ProxyResourceManager 638
 - PSIVTRun script 857, 861
 - Publication consumer 182
 - publications
 - suppressing local ones 950
 - publish exit
 - IBM MQ for IBM i sample program 502
 - IBM MQ for UNIX sample program 502
 - IBM MQ for Windows sample program 502
 - publish/subscribe
 - examples
 - Automated airline gate 201
 - Manual airline gate 201
 - lifecycles 201
 - managing 201
 - writing applications 172
 - publish/subscribe installation verification test 857
 - publish/subscribe messaging 871
 - publisher
 - applications
 - similarity with point-to-point 173
 - types 173
 - writing 173
 - fixed topics 173
 - variable topics 173
 - Publisher
 - application 174, 178
 - fixed topic 174
 - variable topic 178
 - put (sample for IBM MQ for z/OS) 585
 - Put method 728, 768
 - PutApplicationName property 743
 - PutApplicationType property 744
 - PutApplName field 40
 - PutApplType field 40
 - PUTASYNCALLOWED object property 949
 - PutDate field 40
 - PutDateTime property 744
 - PutMsgRecFields field 115
 - PutMsgRecOffset field 115
 - PutMsgRecPtr field 115
 - PutTime field 40
 - putting
 - messages 111
 - one message 121
 - putting a datagram, sample program C++ 799
 - putting messages to a distribution list, example 802
- ## Q
- QSG (queue-sharing group) 4
 - queue
 - application 235
 - authority check on MQOPEN 100
 - closing 100, 110
 - dead-letter 440
 - dynamic, creation of 109
 - exclusive access 107
 - handle 100
 - index 54
 - initiation 235
 - model 109
 - name resolution when remote 110
 - object handle 100
 - opening 100
 - putting messages on
 - C++ 802
 - remote
 - putting messages 118
 - using local definition 102
 - using MQOPEN 110
 - reply-to 39
 - resolving name 102
 - retrieving messages from
 - C++ 802
 - shared access 107
 - triggered 247
 - undelivered message 440

- queue attributes for IBM MQ for z/OS
 - sample application 596
- queue management classes 796
- queue manager
 - alias definition 102
 - configuring for clients 635, 852, 1040
 - connecting to 654
 - Java 1047
 - connecting using MQCONN 92
 - connecting using MQCONNX 94, 99
 - disconnecting 99
 - disconnecting from 654
 - Java 1047
 - ini file
 - authorization service 330
 - location of default 92
 - maximum message length 286
 - operations on
 - Java 1045
 - reply-to 39
 - queue manager groups 297
 - queue manager ini file 330
 - queue managers 294
 - Queue objects
 - creating
 - using a session 945
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - setting properties
 - using a URI 945
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - queue-sharing group (QSG) 4
 - queue-sharing groups 300
 - application programming 266
 - QueueManagerName property 721, 775
 - QueueName property 775
 - queues
 - accessing
 - in Java 1050
 - for MQSeries applications 5
 - queues, accessing
 - classes for .NET 654
 - QueueType property 721
 - quiescing connection
 - MQGET 131
 - quiescing queue manager
 - how applications should react 437
 - MQCONN 93
 - MQOPEN 109
 - MQPUT 114

R

- range of ports, specifying for client connections
 - IBM MQ classes for JMS 991
 - IBM MQ for Java 1049
- read ahead 147
 - JMS 994
- Read method 746
- ReadBoolean method 747
- ReadByte method 747

- ReadDecimal2 method 747
- ReadDecimal4 method 747
- ReadDouble method 748
- ReadDouble4 method 748
- ReadFloat method 749
- reading messages 808
- ReadInt2 method 749
- ReadInt4 method 749
- ReadLong method 749
- ReadNullTerminatedString 750
- ReadShort method 750
- ReadString method 750
- ReadUInt2 method 751
- ReadUnsignedByte method 751
- ReadUTF method 751
- reason codes 778
- ReasonCode property
 - Message class 733
- MQDistributionList class 766
- MQDistributionListItem class 775
- MQGetMessageOptions class 762
- MQPutMessageOptions class 759
- MQQueue class 721
- MQQueueManager class 704
- MQSession class 692
- ReasonCodeName method 694
- ReasonName property
 - MQDistributionList class 767
 - MQDistributionListItem class 776
 - MQGetMessageOptions class 762
 - MQMessage class 733
 - MQPutMessageOptions class 759
 - MQQueue class 722
 - MQQueueManager class 705
 - MQSession class 692
- receive exit
 - program 358
- receiving
 - messages 950
 - poison messages 950
- reconnectable clients 662
- RecordFields property 760
- recoverable resource manager services (RRS)
 - batch adapter 261
- Recovering Transactions 642
- Recovering transactions: Use cases 646
- RecsPresent field 114
- reference guide for IBM MQ automation
 - classes for ActiveX 688
- reference messages 157
- reference-message header
 - message exit program 360
- refreshing the OAM after changing a user's authorization 326
- remote queue
 - using local definition of 102
 - using MQOPEN 110
- RemoteEvent property 705
- RemoteQueueManagerName property 722
- RemoteQueueName property 722
- Remove method 788
 - ItemIndex parameter 788
 - Selector parameter 788
- reopen, implicit operation 805
- reply messages, IMS bridge 68, 280

- reply-to queue 39
- reply-to queue manager 39
- ReplyToQ field 39
- ReplyToQMgr field 39
- ReplyToQueueManagerName property 744
- ReplyToQueueName property 745
- report
 - confirmation of arrival (COA) 11
 - confirmation of delivery (COD) 11
 - exception 11
 - expiry 11
 - negative action notification (NAN) 11
 - positive action notification (PAN) 11
- report message
 - creating 438
- report options, message 1005
- Report property 745
- reports
 - segmented messages 13
- ResizeBuffer method 752
- resolution of queue names 102
- ResolvedQMGrName field 114
- ResolvedQName field
 - MQGMO structure 132
 - MQPMO structure 114
- ResolvedQueueManagerName property 722, 760
- ResolvedQueueName property 723, 760, 763
- resolving local queue names
 - MQOPEN 109
- resource adapter, IBM MQ 1092
- resource enlistment 636
- resource manager, XA compliant name 231
- resources
 - closing using JMS 955
- ResponseRecOffset field 115
- ResponseRecPtr field 115
- restrictions and variations
 - to core classes 1092
- restrictions in z/OS batch 87
- RETAIN_PUBLICATION 995
- retained publications
 - in JMS 995
- RetentionInterval property 723
- retrieving administered objects from JNDI 931
- retrieving items within a message, example 808
- retrieving messages from named queue, example
 - C++ 802
- retry sending message 439
- return codes 257
- ReturnedLength field
 - MQGMO structure 132
- RPG language
 - on IBM i 407
 - support for 455
- RRS
 - calling the stub dynamically with IBM MQ for z/OS 432

- RRS (recoverable resource manager services)
 - batch adapter 261
- run time
 - creating connection factories
 - using the IBM JMS extensions 932
 - using the IBM MQ JMS extensions 939
 - creating destinations
 - using a session 945
 - using the IBM JMS extensions 932
 - using the IBM MQ JMS extensions 939
- runmqdnm 675
- runmqtmc 306
- runmqtmc monitor 250
- runmqtrm monitor
 - how to run 249
- running
 - applications under CICS Transaction Server 1039
 - IBM MQ for Java programs 1086
 - IVT program
 - IBM MQ classes for JMS point-to-point 853
 - IBM MQ classes for JMS publish/subscribe 857
 - IBM MQ resource adapter 1135
 - JMS Samples 1135
 - running a program automatically
 - an example 514
 - how triggering works 234
 - running applications under z/OS UNIX System Services 823
 - running client applications
 - CICS 308
 - Microsoft Transaction Server 42, 310
 - MQ JMS 311
 - Tuxedo 308
 - WebSphere Application Server 311
 - running samples on z/OS 821
 - Running the ActiveX Starter samples
 - MQAX Starter sample for Microsoft Excel 95 or later 795
 - Running the Bank demonstration 795
 - Running the simple sample 795
 - Running the HTML sample 795
 - starter sample using an ActiveX compatible WWW browser 795
 - starter samples for Visual Basic
 - Version 4 or later
 - MQAXDLST sample 794
 - Running the sample 792
 - Starting the MQAXCLSS sample 794

S

- sample 550
- sample application
 - using to verify 1040
- sample applications
 - API-crossing exit for z/OS 263
- Sample applications 633

- sample applications for IBM MQ for z/OS
 - browse 590
 - compiling 576
 - credit check 604
 - features of MQI demonstrated 569
 - get 587
 - logging on to CICS 581
 - mail manager 597
 - message handler 617
 - preparing in batch 576
 - preparing in CICS Transaction Server for OS/390 581
 - preparing in IMS 584
 - preparing in TSO 579
 - print message 592
 - put 585
 - queue attributes 596
- sample client programs 309
- sample cluster workload exit 380
- sample code
 - classes for .NET 651
- sample programs 791
 - DPUT (imqput) 802
 - HELLO WORLD (imqwrl) 798
 - preparing and running 470, 543, 574
 - Compaq NonStop Kernel 473, 474
 - IBM i 472
 - Windows 473, 474
 - SGET (imqset) 802
 - SPUT (imqsp) 802
- Sample programs
 - AMQSSSLC 520
- sample programs for IBM i
 - AMQ0ECHA 491
 - AMQ0REQ4 512
 - AMQSECHA 491
 - AMQSPUT4 503
 - AMQSREQ4 512
 - put 503
 - using triggering 516
- sample programs for IBM MQ for IBM i
 - AMQ0GET4 492
 - AMQ0SETA 518
 - AMQSAPT 477
 - AMQSCLM 540, 547
 - AMQSERV4 522
 - AMQSGBR4 478
 - AMQSGET4 492
 - AMQSSETA 518
 - AMQSTRG4 522
 - AMQSVFC4 482
 - asynchronous put 477
 - MQSTAT 477
 - using remote queues 539
- sample programs for Windows systems and UNIX systems
 - amaspse0 502
 - amq0gbr0 478
 - amq0get0 492
 - amq0put0 503
 - amq0req0 512
 - amqiech2 491
 - amqiechx 491
 - amqiinq2 496
 - amqiinqx 496
 - amqiset2 518

- sample programs for Windows systems and UNIX systems (*continued*)
 - amqisetx 518
 - amqmech2 491
 - amqmechx 491
 - amqminq2 496
 - amqminqx 496
 - amqmset2 518
 - amqmsetx 518
 - amqrgm 505
 - amqrgma 505
 - amqsapt 477
 - amqsbcg 480
 - amqsbcg0 480
 - amqsbcgc 480
 - amqscic0 481
 - amqsclma.c 540, 547
 - amqsdlq 489
 - amqsech 491
 - amqsecha 491
 - amqsechc 491
 - amqsgbr 478
 - amqsgbr0 478
 - amqsgbrc 478
 - amqsget0 492
 - amqsgetc 492
 - amqsinq 496
 - amqsinqa 496
 - amqsinqc 496
 - amqsiqma 497
 - amqsprma 505
 - amqsprma 505
 - amqsptl0 490
 - amqsput0 503
 - amqsputc 503
 - amqsreq 512
 - amqsreq0 512
 - amqsreqc 512
 - amqsset 518
 - amqsseta 518
 - amqssetc 518
 - amqstrg 522
 - amqstrg0 522
 - amqstrgc 522
 - amqstxgx 538
 - amqstxpx 538
 - amqstxsx.c 524
 - amqsvfc0 482
 - amqsvfcx 482
 - amqsxab0.sqb 483
 - amqsxab0.sqc 483
 - amqsxaf0.sqb 483
 - amqsxaf0.sqc 483
 - amqsxag0.c 483
 - amqsxag0.cbl 483
 - amqsxas0.sqb 483
 - amqsxas0.sqc 483
 - amqsxrm 505
 - amqsxrma 505
 - asynchronous put 477
 - browse 478
 - browser 480
 - CICS transaction 481
 - cluster queue monitoring 540, 547
 - data conversion 482
 - dead-letter queue handler 489
 - distribution list 490

- sample programs for Windows systems and UNIX systems *(continued)*
 - echo 491
 - get 492
 - inquire 496
 - inquire properties of a message handle 497
 - MQSTAT 477
 - publish exit 502
 - put 503
 - Reference Messages 505
 - request 512
 - set sample 518
 - trigger monitor 522
 - TUXEDO 524
 - TUXEDO get 538
 - TUXEDO put 538
 - using remote queues 539
 - using triggering 514
 - XA transaction manager 483
- sample security exit 355, 357
- Scope property 723
- scope, handles 93
- scripts provided with IBM MQ classes for JMS 861
- secondary initialization 317
- Secure Sockets Layer
 - certificate revocation list (CRL) 1079
 - CipherSpecs 663, 1078
 - CipherSuites 1078
 - distinguished names 1079
 - distinguished names (DN) 664
 - enabling 663
 - Java 1078
 - handled by JSSE 971, 1077
 - introduction 971
 - Java 1077
 - properties
 - SSLCERTSTORES 973
 - SSLCIPHERSUITE 971
 - SSLFIPSREQUIRED 972
 - SSLPEERNAME 972
 - SSLRESETCOUNT 974
 - renegotiating the secret key 1080
 - sslCertStores field 1080
 - sslCipherSpec property 663
 - sslCipherSuite field 1078
 - sslFipsRequired field 1078
 - sslPeerName field 1079
 - sslPeerName property 664
 - sslResetCount field 1080
 - using JMS 971
- security
 - exit program 349
 - CLNTCONN and SVRCONN channels 357
 - exit program, overview 342
- security enabling interface (SEI) 330
- security exit
 - parameters 355
 - sample 355, 357
- security policy configuration file, editing 845
- Security policy definition file, editing 1037
- Security Services Programming Interface (SSPI) 357
- Segmentation field
 - MQGMO structure 132
- segmented messages 37
 - reports 13
- segmented messages, IMS bridge 69, 281
- SegmentStatus field
 - MQGMO structure 132
- SEI (IBM MQ security enabling interface) 330
- selecting messages 875, 950
- selector for attributes 219
- selector method 789
- Selector method
 - ItemIndex parameter 789
 - OutSelector parameter 789
 - Selector parameter 789
- Selector parameter
 - Count Property 784
 - Item Property 783
- selectors
 - message 875, 950
 - message, and SQL 27
- send
 - exit program 358
- send message, retry on failure 439
- SENDCHECKCOUNT object
 - property 949
- sending
 - messages 949
- sense codes, IMS 65, 278
- serialized applications 267
- server
 - connecting to a client 294
- server bindings connections 632, 1191
- server environment
 - TUXEDO 524
- server session pool 1008
- server-connection channel
 - security exit 357
- service component
 - authorization 330
 - creating your own 329
 - multiple 321
 - stanza 323
- service stanza 322, 328
- ServiceInterval property 723
- ServiceIntervalEvent property 724
- session
 - creating 943
- Session interface 871
- set and inquire 657, 1061
- setjmsenv script 861
- setjmsenv utility 833
- setjmsenv64 script 861
- setMessageLength method 807
- setting
 - connection factory properties summary 930
 - using the IBM JMS extensions 932
 - using the IBM MQ JMS extensions 939
 - destination properties summary 930
 - using a URI 945
- setting *(continued)*
 - destination properties *(continued)*
 - using the IBM JMS extensions 932
 - using the IBM MQ JMS extensions 939
- setting attributes 218
- setting attributes on IBM MQ for z/OS 596
- SGET sample program 802
- Shareability attribute 107
- Shareability property 724
- shared access to a queue 107
- shared queue 4
- shared queues 300
 - application programming 266
 - initiation queue 268
 - SYSTEM.* queues 268
- sharing conversations
 - classes for Java 1069
- Sharing conversations 289
- sharingConversations
 - variable in Java 1069
- SharingConversations
 - MQCD 289
 - MQCXP 289
- signal handling on UNIX products 88
 - threaded applications 89
 - unthreaded applications 88
- Signal1 field 132, 160
- Signal2 field
 - MQGMO structure 132
- signaling 160
- single-phase commit 222
- Single-phase commit 636
- skipping backout 161
- software
 - prerequisites 827
- software, prerequisites
 - classes for .NET 632
 - IBM MQ for Java 1022
- solving problems 635
 - IBM MQ classes for JMS general 863
 - IBM MQ for Java 1041
- SPUT sample program 802
- SQL for message selectors 27, 875
- SQL on IBM i 407
- sslCertStores field 1080
- SSLCERTSTORES object property 973
- sslCipherSpec property 663
- sslCipherSuite field 1078
- SSLCIPHERSUITE object property 971
- sslFipsRequired field 1078
- SSLFIPSREQUIRED object property 972
- sslPeerName field 1079
- SSLPEERNAME object property 972
- sslPeerName property 664
- sslResetCount field 1080
- SSLRESETCOUNT object property 974
- SSPI (Security Services Programming Interface) 357
- STANDARD binding 94
 - environment variable 99
- stanza
 - authorization service 331
 - authorization service, Windows 325

- starter samples 791
- starting
 - connection 942
- starting applications automatically
 - an example 514
 - how triggering works 234
- StartStopEvent property 705
- startup dialog 342
- static XA resource management
 - structure 232
- stream message 875, 1172
- strings, manipulating 806
- StrucId field
 - MQGMO structure 130
 - MQPMO structure 114
- structure
 - MQCXP 359
- structures 79
- stub program for IBM MQ for z/OS
 - CSQBSTUB
 - calling dynamically 427
 - CSQCSTUB
 - calling dynamically 427
 - CSQQSTUB 426
 - IMS 426
- stub programs 80
- styles of messaging 871
- subscriber
 - applications
 - examples 181
 - patterns 181
 - styles 181
 - managed 185, 191
 - message arrival order 208
 - self-managed 185
 - unmanaged 191
- subscription
 - concentration 182
 - control 182
 - durable 182, 185, 191
 - managed 182, 191
 - multiprocessing 182
 - non-durable 185, 191
 - on demand 191
 - unmanaged 191
- Supported Environments 553
- suppressing local publications 950
- switches 814
- sync point
 - calls by platform 78
 - considerations 287
 - coordination 307
- sync point, performance
 - considerations 48
- synchronizing channels 342
- syncpoint
 - considerations 223
 - external manager interfaces 231
 - IMS applications 59, 271
 - in CICS for IBM i applications 227
 - single-phase commit 222
 - two-phase commit 222
 - with IBM MQ for AIX 228
 - with IBM MQ for IBM i 228, 233
 - with IBM MQ for Windows 228
 - X/Open XA interface 231

- syncpoint, performance
 - considerations 55
- SyncPointAvailability property 706
- System.EnterpriseServices
 - .ServicedComponent 639
- SYSTEM.* queues, shared 268
- System.Transactions 639
- System.Transactions namespace 638

T

- TCP/IP
 - client verifying 1040
 - connection, programming
 - Java 1044
- techniques with IBM MQ 50
- temporary dynamic queue
 - closing 110
- text message 875, 1172
- threaded applications, UNIX 89
- threading 687
- time and date of messages 41
- TMI (trigger monitor interface) 250
- tokens, connection pooling 1069
- Tomessage method 790
- ToMessage method
 - Message parameter 790
 - OptionsBag parameter 790
- Topic objects
 - creating
 - using a session 945
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
 - setting properties
 - using a URI 945
 - using the IBM JMS
 - extensions 932
 - using the IBM MQ JMS
 - extensions 939
- topics
 - accessing
 - in Java 1050
- topics, accessing
 - classes for .NET 654
- TotalMessageLength property 745
- trace entries for CICS adapter 433
- trace file name and directory 777
- tracing
 - turning on 1041
- tracing programs
 - .NET 679
- transacted sessions
 - JMS 943
- transaction class 638
- Transaction Processing 643
- transaction processing system 639
- translation of data 165
- transmission header
 - message exit program 360
 - message-retry exit program 362
- transmission of messages
 - maximum transmission size 359
 - transmission buffer 359
- TransmissionQueueName property 724
- TrigData attribute 240

- trigger 234
 - event 235
 - conditions for 241
 - controlling 245
 - message
 - definition 235
 - properties 252
 - without application messages 243
 - monitor
 - for IBM MQ for IBM i 251
 - what it is 236
 - writing your own 250
 - monitor, provided
 - amqltmc0 248
 - AMQSERV4 248
 - AMQSTRG0 248
 - AMQSTRG4 248
 - by platform 248
 - CKTI 248
 - CSQQTRMN 248
 - runmqtmc 248
 - runmqtrm 248
 - process definition 235
 - type of 246
 - trigger messages
 - units of work 247
 - trigger monitor for IBM MQ clients 306
 - trigger monitor interface (TMI) 250
 - TriggerControl attribute 245
 - TriggerData attribute 235
 - TriggerData property 725
 - TriggerDepth attribute 246
 - TriggerDepth property 725
 - triggering 234
 - application design 247
 - how it works 236
 - how it works with the samples 514
 - introduction 234
 - Java applications on IBM MQ for IBM i 251
 - points to note 237
 - prerequisites 239
 - sample program
 - for IBM i 516
 - sample trigger monitor for IBM MQ
 - for Windows 522
 - sample trigger monitor for IBM MQ
 - on UNIX systems 522
 - sequence of events 236
 - setting conditions 245
 - when it does not work 253
 - with the request sample on IBM MQ
 - for Windows 514
 - with the request sample on IBM MQ
 - on UNIX systems 514
 - without application messages 243
 - triggering for IBM MQ for z/OS
 - sample application 608
 - triggering using samples 514
 - TriggerInterval property 706
 - TriggerMessagePriority property 725
 - TriggerMsgPriority attribute 245
 - TriggerType attribute 245
 - TriggerType property 725
 - troubleshooting 776
 - Truncate method 790
 - ItemCount parameter 790

- trusted applications 94, 99
- truststore, making changes
 - using IBM MQ classes for Java 1081
 - using IBM MQ classes for JMS 975
- Tuning nonpersistent messages in AIX 151
- Tuxedo
 - preparing and running client applications 308
 - sample client programs 309
- TUXEDO sample makefile for IBM MQ for Windows 532, 535
- TUXEDO sample programs
 - amqstxgx 524
 - amqstxpx 524
 - amqstxsx 524
 - building server environment 524
- TUXEDO ubbstxcn.cfg example for IBM MQ for Windows 531, 534
- two-phase commit 222, 636
- types of AMQP message 1172
- types of JMS message 875
- types of message body 949
- typical class path settings 833
- typical classpath settings 1029

U

- ubbstxcn.cfg example for IBM MQ for Windows 531, 534
- undelivered message queue, using 440
- undelivered messages, IMS bridge 65, 277
- uniform resource identifiers (URIs) 945
- unit of work
 - message persistence 38
- Universal Resource Identifier 1196
- UNIX products, signal handling 88
- UNIX System Services
 - building application for IBM MQ for z/OS 427
- unmanaged mode 639
- Unmanaged mode 636
- unsupported functions
 - C++ 807
- unthreaded applications, UNIX 88
- URI 1196
 - keywords 1186
 - parameters 1186
 - syntax 1196
- Usage property 726
- use of message types 10
- useEmptyBuffer method 807, 811
- useFullBuffer method 807
- user exit
 - writing and compiling 381
- user exits 263
 - specifying
 - classes for .NET 661
 - using 660
 - with TLS 660
 - writing 660
- user-exit programs 341
 - writing and compiling 344
- UserId property 746
- UserIdentifier field 40

- using
 - IBM MQ for Java 1040
- using C from C++ 805
- using data conversion 685
- using trace 777
- Using WMQDotnetXAMonitor Application 644
- utilities provided with IBM MQ classes for JMS 861

V

- validation
 - of user IDs 360
- Value parameter
 - Count Property 784
 - Item property 784
- verifying
 - TCP/IP clients 1040
 - with the sample application 1040
- Version field
 - MQGMO structure 130
 - MQPMO structure 114
- versions of software required 827
 - IBM MQ for Java 1022
- Visual Basic
 - amqscnxb.vbp sample 450
 - link libraries 293
 - linking 293
 - linking applications 293
 - MQCNOC structure 450
 - MQCONNXAny call 450
 - on Windows NT 417
 - support for 450
- Visual C++ 816
- VisualAge C++ 816

W

- waiting for messages 159
- WaitInterval field 131, 159
- WaitInterval property 763
- WebSphere Application Server 42, 1097
 - preparing and running client applications 311
- when your IBM MQ automation classes for ActiveX script fails 778
- WMQ_MQMD_MESSAGE
 - _CONTEXT 963
- WMQ_MQMD_READ_ENABLED 963
- WMQ_MQMD_WRITE_ENABLED 963
- work header, writing a message to 814
- write method 807
- Write method 752
- WriteBoolean method 753
- WriteByte method 753
- WriteDecimal2 method 753
- WriteDecimal4 method 754
- WriteDouble method 754
- WriteDouble4 method 754
- WriteFloat method 755
- WriteInt2 method 755
- WriteInt4 method 755
- writelnItem method 807
- WriteLong method 755
- WriteNullTerminatedString method 756

- WriteShort method 756
- WriteString method 756
- WriteUInt2 method 757
- WriteUnsignedByte method 757
- WriteUTF method 757
- writing
 - channel exits 660
 - channel exits in Java
 - for IBM MQ classes for Java 1063
 - for IBM MQ classes for JMS 978
 - JMS applications 871
 - programs 647
 - Java 1043
 - user exits 660
- writing a cluster workload user exit 381
- writing applications 283
- writing exit programs
 - data conversion
 - IBM MQ for IBM i 367
 - IBM MQ for z/OS 368
 - UNIX systems 369
 - Windows NT 372
- writing messages
 - to the CICS bridge 813
 - to the dead-letter queue 812
 - to the IMS bridge 812
 - to the work header 814

X

- X/Open XA interface support 231
- XA 636
- XA resource manager
 - name 231
- XA transaction manager samples 483
- xa_commit 638
- xa_end 638
- xa_open 638
- xa_prepare 638
- xa_start 638

Z

- z/OS
 - batch restrictions 87
 - differences with 1089
 - support for 259
 - UNIX System Services 262
- z/OS compiling 821

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[®], 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. 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
- Use the form on the web here: 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.

