Version 8 Release 0

IBM MQ Appliance

Version 8 Release 0

IBM MQ Appliance

Note
Before using this information and the product it supports, read the information in "Notices" on page 627.

This edition applies to version 8 release 0 of IBM MQ Appliance 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 2015, 2016.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Product overview 1

Introduction to the IBM MQ Appliance 1
Appliances and the appliance firmware 2
Appliance specification 2
What's new 3
What's changed 3
High availability 3
Disaster recovery for a high availability configuration 6
Disaster recovery 6
Accessibility features for IBM MQ Appliance . . . 7
Notices 8
Trademarks 11
Terms and conditions for product documentation 11
IBM Online Privacy Statement 12

Chapter 2. Planning 13

Differences between administering an IBM MQ Appliance and an IBM MQ installation 13
Control commands on the IBM MQ Appliance. . 13
Differences between queue managers that are running on the IBM MQ Appliance and an IBM MQ installation 18
Exits and services on the IBM MQ Appliance . . 18
Queue manager configuration on the IBM MQ Appliance 19
IBM MQ objects on the IBM MQ Appliance . . 20
XA transactions 20
Planning a high availability system 21
Network requirements for high availability . . . 22
Planning a disaster recovery system 23
Hardware limitations for disaster recovery . . . 24
Planning network connections 25
Capacity planning 28

Chapter 3. Installing 31

Safety 31
Guidelines for servicing electrical equipment . . 32
Inspecting for unsafe conditions 33
Safety statements 34
Introducing the IBM MQ Appliance 37
Specifications and features 38
Intrusion detection 39
Components on the front 39
Components on the rear 46
Prepare for installation 46
Rack requirements 47
Tool requirements 48
Installing the appliance in a rack 48
Installing rails in the rack frame 48
Installing the appliance on the rails 50
Considerations to connect the appliance to an AC power source 52
Connect the appliance to a network 53
Setting up the initial firmware configuration . . . 54
Configuration requirements 54
Firmware considerations 55
Procedure 1 of 3: Connecting the serial cable to the appliance 56
Procedure 2 of 3: Initializing the appliance . . . 57
Procedure 3 of 3: Accepting the license agreement 58
Maintenance 59
Diagnosing your appliance 59
Troubleshooting your appliance 65
Removing or replacing the appliance or parts . . 68
Removing the batteries 89

Chapter 4. Upgrading and downgrading 93

Installing a new level of firmware 93
Reverting to the previous level of firmware. . . . 94
Downgrading. 95
Suspending an appliance from an HA group for maintenance 96

Chapter 5. Configuring 97

Command line access 97
Connecting to the serial port 97
Configuring the SSH service 97
Configuring the IBM MQ Appliance web UI . . . 99
Changing the IBM MQ Appliance web UI IP address and port 99
Configuring certificates for IBM MQ Appliance web UI 100
Configuring the appliance 103
Ethernet interfaces. 103
VLAN interfaces 106
Link aggregation interfaces 109
IPMI Settings 111
DNS settings 112
Configuring the locale, date, and time 114
Configuring the appliance name 118
Configuring appliance users 120
Configuring queue managers 123
Configuring environment variables 124
Adding an environment variable 124
Modifying an environment variable 125
Removing an environment variable 126
Viewing environment variables 127
Configuring high availability 127
Configuring the hardware for high availability 128
Configuring the high availability group . . . 130
Configuring high availability queue managers 134
Example network set up for HA configuration 138
Configuring disaster recovery for a high availability queue manager 140
Replacing a failed node in a high availability group 143
Configuring disaster recovery 144
Configuring the hardware for disaster recovery 144
Configuring disaster recovery queue managers 146

© Copyright IBM Corp. 2015, 2016

Configuring disaster recovery for a high availability queue manager 152

Chapter 6. Administering 157

Using the IBM MQ Console 157
Working with queue managers 157
Working with IBM MQ objects 160
Working with authority records 175
Monitoring system resource usage by using the IBM MQ Console 178
Importing and exporting dashboard layouts . . 189
Message queue control commands 190
Administering messaging users 190
Using MQSC commands 192
Using an IBM MQ client. 192
Setting up a queue manager to accept client connections 192
Configuring queue managers and objects by using a client 194
Putting and getting messages 195
Publishing and subscribing 195
Browsing a message queue 196
Creating and downloading a CCDT file 197
Starting and stopping the appliance 197
Restarting the appliance 198
Shutting down the appliance 198
Back up and restore 198
Backing up or saving the appliance configuration 199
Restoring the appliance configuration 199
Backing up messaging users 200
Restoring messaging users 201
Backing up a key repository 202
Restoring a key repository 203
Backing up a queue manager 204
Restoring a queue manager. 205
Backing up IBM MQ Appliance web UI configuration data 206
Restoring IBM MQ Appliance web UI configuration data 207
Backing up and restoring message data . . . 208
Factory reset. 208
Operating in a high availability environment . . . 208
Suspending an appliance from an HA group for maintenance. 209
Replacing a failed node in a high availability group 210
Managing queue manager locations in a high availability group 211
Viewing the status of appliances in a high availability group 213
Viewing the status of a high availability queue manager 214
Disaster recovery for a high availability queue manager 215
Operating in a disaster recovery environment . . 215
Switching over to a recovery appliance 216
Switching back to the main appliance 217
Replacing a failed node in a disaster recovery configuration 227
Testing the recovery appliance. 230
Reversing disaster recovery roles 230
Viewing the status of a disaster recovery queue manager 231
Managing files by using the IBM MQ Appliance web UI 233
Watchdog timer 236

Chapter 7. Migrating and consolidating 239

Moving queue managers from other IBM MQ platforms. 239
Moving a queue manager 240
Moving queue managers secured by using TLS 242
Planning for incompatible features in the queue manager 243
Handling incompatible features in the queue manager 244
Editing qm.ini files 245
Transferring queue managers to other IBM MQ Appliances 249
Transfer from an existing single appliance to a new single appliance 249
Transfer from an existing appliance in a disaster recovery configuration 253
Transfer from an existing high availability pair of appliances to a new pair of appliances . . . 258

Chapter 8. Security 265

User authentication and authorization 265
TLS certificate management 265
Working with self-signed certificates. 266
Working with CA-signed certificates. 268
Listing certificates for a queue manager . . . 274
Viewing a certificate for a queue manager . . . 275
Deleting a certificate 275
Managing certificates on the appliance 276

Chapter 9. Monitoring and reporting 283

Monitoring system resource usage 283
Monitoring system resource usage by using the status command 283
Monitoring system resource usage by using the amqsrua command 284
Monitoring system resource usage by using the IBM MQ Console 286
Monitoring the appliance by using the show command 296
Developing your own resource monitoring program 298
Interpreting CPU monitoring information . . . 299
Application activity trace 300
Subscriptions to application activity trace . . . 300
Creating subscriptions to application activity trace 301
Application activity trace: subscriptions compared with central collection 302
Using **amqsact** to view trace messages 303
Configuring trace levels 305
System topics for monitoring and activity trace . . 306

Chapter 10. Troubleshooting 309
Error logs 309
Viewing system error log files 310
Viewing queue manager error log files 311
Viewing the first failure data captures 312
Deleting log files 313
Downloading error logs 313
Reason codes 314
Event logs 315
Types of log target. 315
Configuring log targets 316
Using trace 316
Using trace in the IBM MQ Console 317
Resolving a partitioned problem in a high availability configuration 318
Resolving a partitioned problem in a disaster recovery configuration 319
Troubleshooting file copy 321
Help with using **runmqras** 322

Chapter 11. Reference. 325
IBM MQ commands 325
IBM MQ Control commands 325
IBM MQ configuration commands 365
Messaging user and group commands 370
Queue manager certificate management commands 374
Queue manager commands. 389
High availability commands 408
Troubleshooting commands 415
Appliance commands 436
Appliance commands 436
Appliance user commands 449
Appliance user group commands. 451
Crypto commands. 454
Date, time, and locale configuration commands 478
DNS commands 490
Ethernet commands 495
File commands 505
Global configuration commands 509
Host Alias commands 530
Initialization commands 531
IPMI user commands. 536
IPMI LAN channel commands 538
Link aggregation commands 542
Log target commands 551
Monitoring and reporting commands 572
Network commands 599
SSH client commands 605
System Settings commands 610
VLAN Commands. 614
Web management service commands 623
Messages 626

Notices 627
Programming interface information 629
Trademarks 629

Chapter 1. Product overview

The IBM® MQ Appliance is an appliance-based offering of IBM MQ.

For an introduction to and overview of IBM MQ, see IBM MQ Technical overview in the IBM MQ documentation.

Introduction to the IBM MQ Appliance

The IBM MQ Appliance provides IBM MQ V8 on an appliance. You can create one or more queue managers on an appliance and connect them as part of an IBM MQ network.

The IBM MQ Appliance is designed to be easy to deploy. It has a command-line interface and a web UI for configuring and administering the appliance, and a web-based user interface, the IBM MQ Console, for administering queue managers. You can set up multiple administration user accounts on the appliance.

The IBM MQ Appliance can provide a number of messaging solutions in your enterprise:

Rationalize your existing installation
If you are an existing IBM MQ user, you might want to use appliances to consolidate your IBM MQ architecture. You might be in the situation where you have a large messaging estate that is spread across a number of different hardware platforms. Maintaining and keeping this estate up to date can prove to be a large overhead; simplifying your estate by adding IBM MQ Appliances can greatly reduce your total cost of operation.

Install an easy-to-update solution from the start
If you are a new IBM MQ user, using IBM MQ Appliances from the outset can make your solution easier to later extend. You just add more appliances.

Implement high availability
If you need 24/7 reliability, you can pair appliances up to provide a high availability solution. If your system fails, you are rapidly switched over to a replica of that system.

Implement disaster recovery
You can implement disaster recovery to manually switch over to a replica in a different data center should your system fail.

Place a messaging server in outlying premises
If you have outlying offices, factories, or branches, you can place an IBM MQ Appliance on those premises to provide a simple, distributed messaging solution. You could also employ the same solution for a business partner, by putting an appliance on their premises.

The IBM MQ Appliance can be easily updated by downloading and installing a new firmware version. You can install certificates onto the appliance, and connect to your IBM MQ V8 network by using TLS.

You run applications on clients that connect to the appliance.

© Copyright IBM Corp. 2015, 2016

View the video introductions to IBM MQ Appliance:

IBM MQ Appliance

Key Features of the IBM MQ Appliance

Appliances and the appliance firmware

The IBM MQ Appliance is a hardware product that provides IBM MQ ready installed and ready to use. There is no general-purpose operating system that is exposed to the administrator or messaging user, and everything that runs on the appliance is factory-installed in the appliance firmware.

The appliance firmware can be updated by downloading update files from IBM. No other installation or maintenance application on the system is required.

You use simple administration interfaces on the appliance to perform the traditional management tasks on your queue managers and the objects that they host. You initially configure the appliance by using a command line interface or the IBM MQ Appliance web UI, see "Configuring the appliance" on page 103.

You can use the IBM MQ Console part of the IBM MQ Appliance web UI for simple browser-based management of your queue managers, see "Using the IBM MQ Console" on page 157. After you have created a queue manager on the appliance, you can use IBM MQ management tools such as the IBM MQ Explorer or **runmqsc** from a remote system, or other IBM or third-party management products, to work with the queue manager.

Because only IBM certified firmware updates can be installed on the appliance, all applications that connect to appliance queue managers do so by using the IBM MQ Client protocol. For information about migrating queue managers and applications on to the appliance platform see Chapter 7, "Migrating and consolidating," on page 239.

A major functional difference in the IBM MQ Appliance product when compared to IBM MQ software is the high availability data replication feature, see "High availability" on page 3, and the disaster recovery feature, see "Disaster recovery" on page 6.

Appliance specification

The table below provides the technical details of the IBM MQ Appliance hardware configuration.

You can use these details to help in performance planning and sizing activities. Do not, however, assume 'like for like' mappings for other installations of IBM MQ, for example, installations on UNIX platforms. For detailed planning, the best source of information when assessing expected performance (message throughput) are the following performance reports:

- IBM MQ Appliance Performance Report
- IBM MQ Appliance HA/DR Performance Report

Table 1. Appliance specification

M2001 (current model)	M2000 (previous model)
2 x 10 core "Ivy Bridge" x86 processors (6 cores active in B model). 2.80 GHz, Hyperthreading enabled.	2 x 10 core "Ivy Bridge" x86 processors (6 cores active in B model). 2.80 GHz, Hyperthreading enabled.
192 GB RAM, 1600 MHz DIMMs	192 GB RAM, 1600 MHz DIMMs
2 x Management 1 Gb Ethernet ports (one supporting IPMI)	2 x Management 1 Gb Ethernet ports (one supporting IPMI)
8 x 1 Gb Ethernet ports	8 x 1 Gb Ethernet ports
4 x 10 Gb Ethernet ports	2 x 10 Gb Ethernet ports
2 x 3.2 TB SSDs under hardware controlled mirrored RAID (3 TB effective)	2x 1.2 TB HDDs under hardware controlled mirrored RAID (1 TB effective)
1 GB RAID cache	1 GB RAID cache

What's new

This topic describes new features in version 8.0.0.5 of the appliance firmware.

The following features are new for version 8.0.0.5:

- Disaster recovery is now available for high availability queue managers. See "Disaster recovery for a high availability configuration" on page 6.
- The appliance now supports AMQP channels. See Developing MQ Light applications in the IBM MQ documentation
- Support for M2001 appliance features (solid-state disk drives and 4 port 10 Gb Ethernet module).

What's changed

This topic describes significant features that have changed in version 8.0.0.5 of the appliance firmware.

- The command prompt has changed with version 8.0.0.5. For either the M2001 or the M2000 models of the appliance the command prompt is:

  ```
  mqa#
  ```

High availability

The IBM MQ Appliance might experience outages both planned and unplanned. The high availability (HA) features of the appliance enable queue managers to have maximum availability. The HA features of the IBM MQ Appliance give the appliance an ability to withstand software or hardware outages. Therefore, it is available as much of the time as possible. These outages might be planned events, such as maintenance and backups, or unplanned events, such as hardware failures or power failures.

To configure HA for the IBM MQ Appliance, you can connect a pair of appliances either directly, or by using switches (a separate switch for each link). You must then create an HA group for this pair of appliances. To work most effectively as a high availability solution, the two appliances need to be in close physical proximity to one another. For this reason the HA solution is not intended to provide disaster recovery , although you can configure a disaster recovery solution for queue managers that run on your HA pair.

For details of combining HA and disaster recovery on the appliance, see “Disaster recovery for a high availability configuration” on page 6. A queue manager can belong to an HA group and be part of a disaster recovery configuration.

The HA group controls the availability of any HA queue managers that are created on the appliances. By default, the HA queue managers are run on the appliance on which they are created, when that appliance is available. This appliance is known as the *preferred appliance* of the HA queue manager. You can use commands to specify the other appliance as the preferred appliance, if required, or to specify that the queue manager has no preferred appliance.

If an appliance in the pair is stopped for any reason, the HA queue managers that are running on that appliance automatically start to run on the other appliance. That is, the queue managers are failed over to the other appliance. When the stopped appliance is restarted, and the required data is replicated back to that appliance, it resumes running the HA queue managers for which it is the preferred appliance. Persistent messages are preserved.

To ensure that the HA queue manager is ready to run on either appliance, queue manager data is replicated synchronously between the appliances. In some situations, such as when one appliance is unavailable, the queue manager data cannot be replicated synchronously. When the appliance becomes available, the queue managers in the HA group enter a catch-up phase, in which the queue manager data is replicated. The appliances use a dedicated 10 Gb Ethernet connection for replication.

This HA solution enables all the HA queue managers in the HA group to continue running when one appliance in the group is stopped. If both appliances in the HA group fail at the same time, the HA queue managers cannot run until at least one of the appliances is restarted.

Appliances in an HA group can run other queue managers that are outside of the HA group, but if the appliance fails or is stopped, then that queue manager stops. Appliances can belong only to one HA group.

For more information about configuring HA on the IBM MQ Appliance, see “Configuring high availability” on page 127. For information about planning the physical location, see “Planning a high availability system” on page 21.

Example HA group

In the example configuration, two IBM MQ Appliances, named castor and pollux, are located in the same data center, in adjacent racks. The three cables that connect the two appliances are less than a meter long, and so communication between the two has the minimum of delay.

The appliance that is named pollux runs one queue manager, terentia1, which is inside the HA group. The appliance that is named castor has two queue managers that are running within the HA group, cicero1 and cicero2. It also has another queue manager, tullia2, that runs outside the HA group. Both castor and pollux have shadow versions of the HA queue managers on the other appliance. These queue managers are kept up to date by replication across the replication interface. Two more interfaces, a primary and a secondary, track the heartbeat of the other appliance.

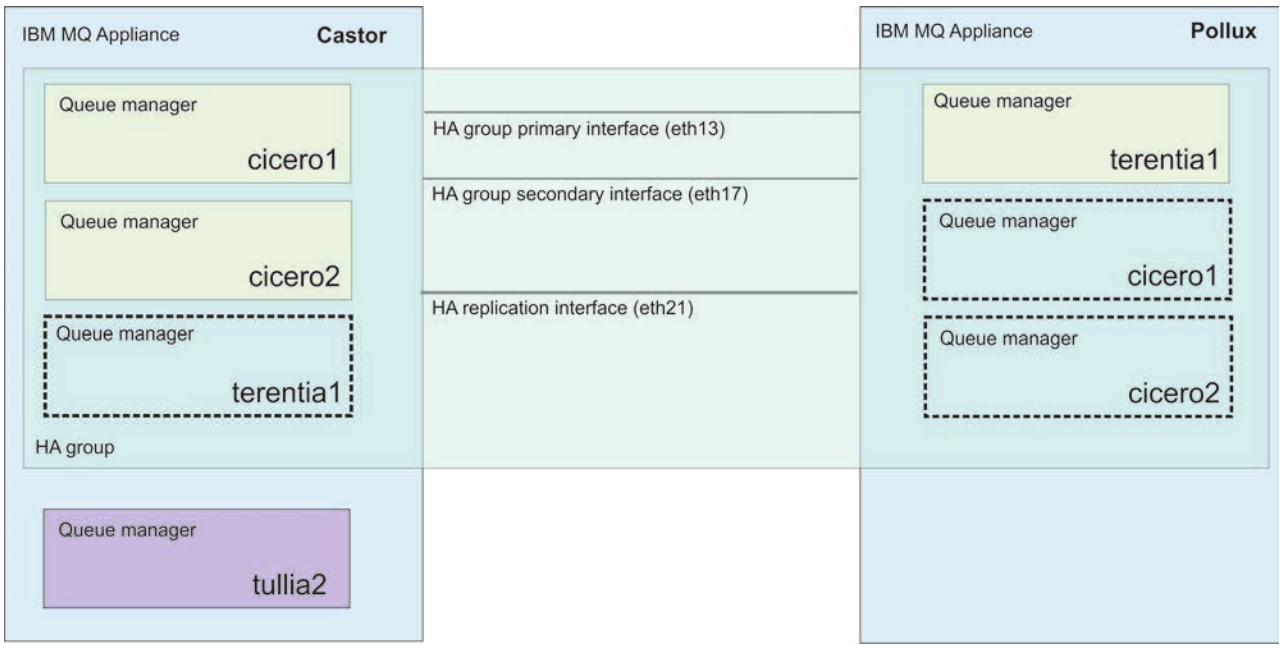

Figure 1. Example HA group

The rack that castor is in suffers a power failure. The appliance that is named pollux detects that castor has failed, and starts to run the queue managers cicero1 and cicero2. The queue manager tullia2 is outside the HA group, so does not fail over to pollux.

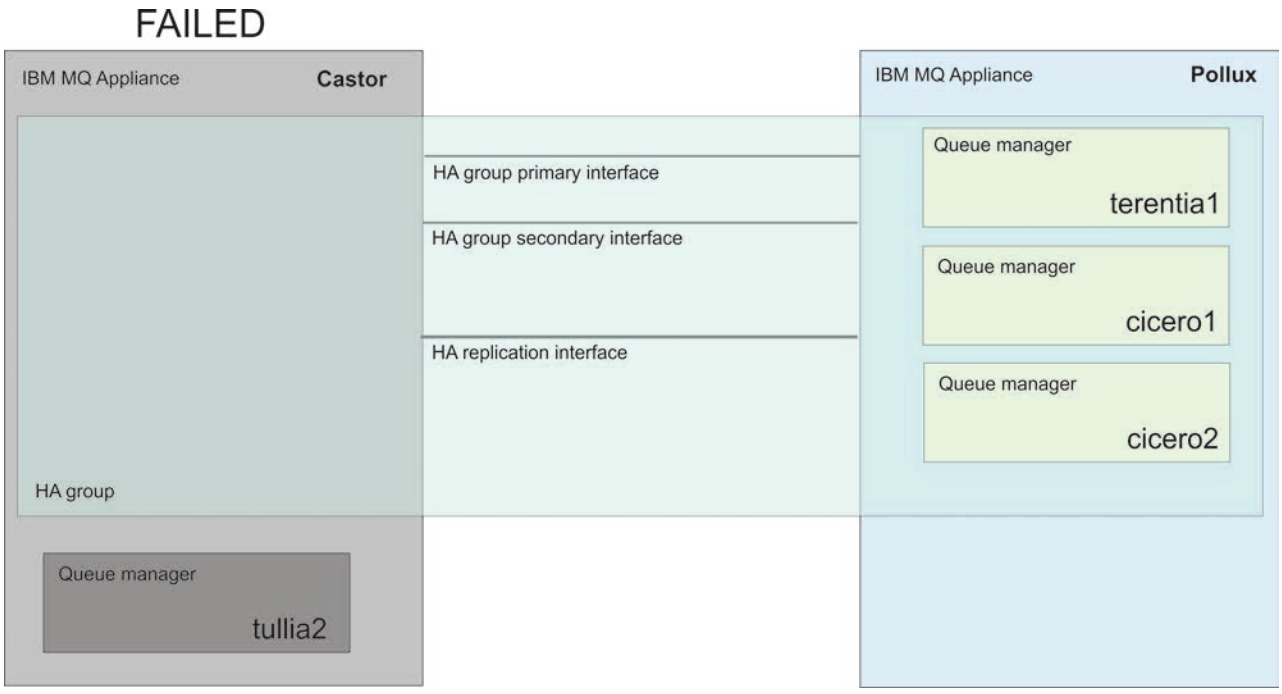

Figure 2. Example HA group after appliance failure

When power is restored, the queue managers cicero1 and cicero2 run on the appliance that is named castor again.

Demonstration HA with put and get operation

View the video for a demonstration of the simple put and get demonstration running on an HA pair of IBM MQ Appliances.

MQ Appliance demo - Simple MQ HA put/get sample

Disaster recovery for a high availability configuration

You can configure a disaster recovery (DR) solution such that an appliance at a remote location can take over if both the appliances in a high availability (HA) pair fail at the same time.

If you are running a queue manager on an appliance in an HA pair, and both appliances in the pair become unavailable (for example, the data center has a major power failure), then you can manually start and run the queue manager on a DR appliance at a different site.

Replication between the appliances in a high availability pair is synchronous, but queue manager data is replicated to the DR appliance asynchronously. This means that following a recovery situation, some messaging data might be lost. But the queue manager on the DR appliance will be in a consistent state, and able to start running immediately, even if it is started at a slightly earlier part of the message stream.

A floating IP address is allocated to each queue manager under HA/DR control. The DR appliance can connect to a single IP address regardless of which of the appliances in the HA group is running the queue manager.

The following diagram shows an HA configuration with disaster recovery.

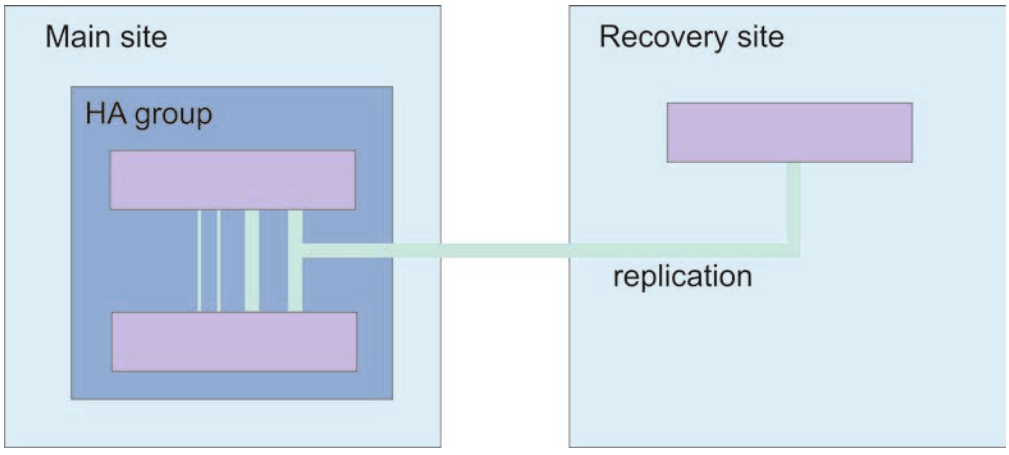

Disaster recovery

The IBM MQ Appliance disaster recovery solution provides for the situation where you have a complete outage at your data center. The work can be resumed by another IBM MQ Appliance running at a distant location.

Disaster recovery (DR) is provided on a per-queue manager basis. When you create a queue manager on your main appliance, you create a secondary instance on your recovery appliance at the distant site. The two appliances are linked by a high-speed connection. The work of the primary queue manager is replicated to the secondary queue manager asynchronously. For example, an IBM MQ PUT or GET completes and returns to the application before the event is replicated to the secondary queue manager. Asynchronous replication means that, following a recovery situation, some messaging data might be lost. But the secondary queue manager will be in a consistent state, and able to start running immediately, even if it is started at a slightly earlier part of the message stream.

You can configure a queue manager so that it is part of a disaster recovery configuration and a high availability group, see "Disaster recovery for a high availability configuration" on page 6.

The main and recovery appliances are connected by a single replication link. Unlike the high availability solution, there is no heartbeat detection between the two appliances. An appliance at the recovery site can host secondary queue managers from multiple appliances at the main site, or at different main sites. For example, you could have an appliance in Glasgow that provided disaster recovery for appliances in Birmingham, Paris, and Frankfurt. Equally, an appliance at your main site could have secondary queue managers on different appliances at different recovery sites.

When a disaster occurs, and the main appliance is lost and a primary queue manager stops running, the secondary queue manager at the distant site can be started manually. Applications must connect to the recovery appliance (using automatic client reconnection). The secondary queue manager can then process application messages until such time as normal operation can be resumed. There can be up to 4 MB of data in the TCP send buffer of a primary queue manager, ready to be replicated to the secondary instance, and this data is lost if a disaster occurs.

Replication, synchronization, and snapshots

When the two appliances in a DR configuration are connected, any updates to the persistent data for a DR queue manager are transferred from the primary instance of the queue manager to the secondary instance. This is known as **replication**.

If the network connection between the appliances is lost, the changes to the persistent data for the primary instance of a queue manager are tracked. When the network connection is restored, a different process is used to get the secondary instance up to speed as quickly as possible. This is known as **synchronization**.

While synchronization is in progress, the data on the secondary instance is in an inconsistent state. A **snapshot** of the state of the secondary queue manager data is taken. If a failure of the main appliance or the network connection occurs during synchronization, the secondary instance reverts to this snapshot and the queue manager can be started. Any of the updates that happened since the original network failure are lost, however.

Accessibility features for IBM MQ Appliance

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Accessibility features

IBM MQ Appliance includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

IBM MQ Appliance uses the latest W3C Standard, WAI-ARIA 1.0, to ensure compliance to US Section 508, and Web Content Accessibility Guidelines (WCAG) 2.0. To take advantage of accessibility features, use the latest release of your screen reader in combination with the latest web browser that is supported by this product.

The IBM MQ Appliance online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described at http://www.ibm.com/support/knowledgecenter/about/releasenotes.html.

Keyboard navigation

This product uses standard navigation keys.

Interface information

The fully accessible way of using IBM MQ Appliance is to use the command line interface. For more information about using commands, see Command reference.

The IBM MQ Appliance user interfaces do not have content that flashes 2 - 55 times per second.

The IBM MQ Appliance web user interface does not rely on cascading style sheets to render content properly and to provide a usable experience. However, the product documentation does rely on cascading style sheets. IBM MQ Appliance provides an equivalent way for low-vision users to use a user's system display settings, including high-contrast mode. You can control font size by using the device or browser settings.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has established a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service 800-IBM-3383 (800-426-3383) (within North America)

IBM and accessibility

For more information about the commitment that IBM has to accessibility, see IBM Accessibility.

Notices

This information was developed for products and services that are offered in the USA.

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, MD-NC119
Armonk, NY 10504-1785
United States of America

For license inquiries regarding double-byte character set (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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
USA

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 document 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 1993,2017. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

The Oracle Outside In Technology included herein is subject to a restricted use license and can only be used in conjunction with this application.

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these

publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

Depending upon the configurations deployed, this Software Offering may use session and persistent cookies for purposes of session management, authentication, or other functional purposes. These cookies cannot be disabled.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see: (i) IBM's Privacy Policy at http://www.ibm.com/privacy; (ii) IBM's Online Privacy Statement at http://www.ibm.com/privacy/details (in particular the section entitled "Cookies, Web Beacons and Other Technologies"); and (iii) the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Chapter 2. Planning

When you plan your IBM MQ Appliance solution, you must ensure that you understand the key differences between the IBM MQ Appliance and IBM MQ.

Planning any deployment of IBM MQ, and applications that communicate by using the messaging services that it provides, requires an understanding of both general messaging concepts and IBM MQ-specific terminology and environment considerations. For IBM MQ general information and information on planning, see IBM MQ Technical overview and Planning in the IBM MQ documentation

These topics describe additional planning considerations for an IBM MQ Appliance deployment, as opposed to an IBM MQ software installation.

Differences between administering an IBM MQ Appliance and an IBM MQ installation

The IBM MQ Appliance provides an environment to configure and manage the resources that are required for your IBM MQ system.

Many IBM MQ administrative concepts and commands are supported on the appliance, although some differences do exist.

Control commands on the IBM MQ Appliance

You can use the IBM MQ control commands on the IBM MQ Appliance command line. However, not all of the control commands are supported and some of the control commands have different parameters to the IBM MQ equivalent.

To use the IBM MQ control commands, you must enter the IBM MQ administration mode by entering the command **mqcli** on the command line. You can exit the IBM MQ administration mode by entering the command **exit**.

The following example shows how to enter the IBM MQ administration mode and create a queue manager:

```
mqa# mqcli
mqa(mqcli)# crtmqm QM1
MQ Appliance queue manager created. Creating or replacing default objects for queue manager 'QM1'.
Default objects statistics : 79 created. 0 replaced. 0 failed. Completing setup.
Setup completed.
mqa(mqcli)# exit
mqa#
```

Unsupported commands

The commands that are not supported are listed in the following table:

Table 2. Unsupported commands

Command	Comment
crtmqcvx	Use of exits is not supported on the appliance. See “Exits and services on the IBM MQ Appliance” on page 18.
crtmqenv	Use the mqcli command line environment “setmqvar” on page 369 command.

© Copyright IBM Corp. 2015, 2016

Table 2. Unsupported commands (continued)

Command	Comment
dltmqinst	See Chapter 4, “Upgrading and downgrading,” on page 93.
dmpmqaut	Replaced by IBM MQ Console (see “Working with authority records” on page 175) or authority record configuration using `runmqsc` (for example, SET AUTHREC) see “runmqsc” on page 354.
dmpmqmsg	Use from an external system, connecting as a client.
dspmqaut	Replaced by IBM MQ Console (see “Working with authority records” on page 175) or authority record configuration using `runmqsc` (for example, SET AUTHREC) see “runmqsc” on page 354.
dspmqcsv	External control of some queue manager components is not supported on the IBM MQ Appliance.
dspmqfls	IBM MQ file access is not supported on the appliance.
dspmqinst	See Chapter 4, “Upgrading and downgrading,” on page 93
dspmqspl	Replaced by the SET/DISPLAY POLICY configuration using `runmqsc`, see “runmqsc” on page 354.
dspmqtrc	Appliance side tracing is for IBM use only.
endmqcsv	External control of some queue manager components is not supported on the IBM MQ Appliance.
migmbbrk	Migration tools for versions of IBM MQ earlier than V8 are not applicable to IBM MQ Appliance queue managers.
rcdmqimg	Linear logging is not supported on the IBM MQ Appliance.
runmqchi	External control of some queue manager components is not supported on the IBM MQ Appliance.
runmqchl	External control of some queue manager components is not supported on the IBM MQ Appliance.
runmqdlq	Use a dead letter queue handler connected as a client from a remote system. As an example, the amqsdlq sample can be compiled and linked as a client application to provide the same behavior as the runmqdlq program. See Dead-letter queue handler sample.
runmqlsr	Listeners must be started by way of queue manager administration. You can use the IBM MQ Console, see “Working with listeners” on page 170, or use `runmqsc`, see “runmqsc” on page 354.
runmqtmc, runmqtrm	Use **`runmqtmc`** (the client form of this command) from a remote MQ server or client installation, on which the triggered application runs.
setmqaut	Replaced by IBM MQ Console (see “Working with authority records” on page 175) or authority record configuration using `runmqsc` (for example, SET AUTHREC) see “runmqsc” on page 354.
setmqenv	Use the mqcli command line environment “setmqvar” on page 369 command.
setmqinst	See Chapter 4, “Upgrading and downgrading,” on page 93
setmqm	See Chapter 4, “Upgrading and downgrading,” on page 93.
setmqspl	Replaced by the SET/DISPLAY POLICY configuration using `runmqsc`, see “runmqsc” on page 354.

Table 2. Unsupported commands (continued)

Command	Comment
setmqprd	See Chapter 4, "Upgrading and downgrading," on page 93
strmqcfg	Run IBM MQ Explorer from an external system.
strmqcsv	External control of some queue manager components is not supported on the IBM MQ Appliance.

For authorization commands, command server commands, and security policy commands, you can use the equivalent PCF or MQSC commands. For more information about the equivalent commands, see Comparing command sets in the IBM MQ documentation.

Supported commands

The IBM MQ commands that you can use on the IBM MQ Appliance are listed in the following table.

Table 3. Supported IBM MQ commands

Command	Comments
"crtmqm" on page 326	Create queue manager. The following parameters are not supported: • `-lc` • `-ll` • `-ld` • `-g` • `-md` • `-oa group` • `-q` • `-si` • `-ss` • `-z` The following parameters have been added: • `-fs` *`FileSize`* • `-sx`
"dltmqm" on page 329	Delete queue manager. The following parameter is not supported: • `-z`
"dmpmqcfg" on page 330	Dump queue manager configuration
"dspmq" on page 333	Display queue managers. The following parameters are not supported: • `-o installation` • `-o standby` • `-x` The following parameters have been added: • `-o ha` • `-o dr`
"dspmqrte" on page 334	Display route
"dspmqtrn" on page 341	Display transactions

Table 3. Supported IBM MQ commands (continued)

Command	Comments
"dspmqver" on page 343	Display version and build information. The following parameter is not supported: • `-i` The `-p` parameter supports only the values 1, 64, and 128. The output of this command is not the same as for the IBM MQ **`dspmqver`** command. Information about the operating system, installation details, and data paths are not displayed. That is, only the name, version, level, and build type information is displayed.
"endmqm" on page 344	End queue manager. The following parameters are not supported: • `-s` • `-x` • `-z`
"endmqtrc" on page 346	End trace
"mqrc" on page 347	IBM MQ return code
"rcrmqobj" on page 349	Generate a client channel definition table (CCDT)
"rsvmqtrn" on page 349	Resolve transaction
"runmqras" on page 351	Run diagnostics collection. The following parameters are not supported: • `-outputdir` • `-zipfile` • `-workdirectory`
"runmqsc" on page 354	Run MQSC commands. The following parameters are not supported: • `-n` • `-c`
"runswchl" on page 356	Switch cluster channel
"strmqm" on page 357	Start queue manager. The following parameters are not supported: • `-x` • `-z` • `-a` • `-r`
"strmqtrc" on page 359	Start trace

New commands

New IBM MQ commands that are specific to the IBM MQ Appliance are listed in the following table:

Table 4. Appliance commands

Command	Description
"dspmqerr" on page 416	Display the IBM MQ error log files.
"crthagrp" on page 408	Create a high availability (HA) group of appliances.

Table 4. Appliance commands (continued)

Command	Description
“dsphagrp” on page 409	Display the status of the appliances in the high availability (HA) group.
“makehaprimary” on page 410	Specifies that an appliance is the 'winner' when resolving a partitioned situation in the high availability group.
“prepareha” on page 411	Prepare an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.
“sethagrp” on page 412	Pause and resume an appliance in a high availability group.
crtdrprimary	Augment an existing queue manager to become the primary queue manager in a disaster recovery configuration.
crtdrsecondary	Create a secondary version of a queue manager on the recovery appliance in a disaster recovery configuration.
makedrprimary	Switch a disaster recovery queue manager to have the primary role in the disaster recovery configuration.
makedrsecondary	Prevent a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.
dltdrprimary	Remove a queue manager currently in the primary role from DR control.
dltdrsecondary	Remove a queue manager currently in the secondary role from DR control and delete it.
“dspmqini” on page 365	Display attributes from the `qm.ini` or `mqat.ini` file of a specified queue manager.
“dspmqvar” on page 366	Display environment variables set for a specified queue manager.
“setmqini” on page 367	Add or remove an attribute from the `qm.ini` file of a specified queue manager. Set a value for an attribute in the `mqat.ini` file.
“setmqvar” on page 369	Add or remove an environment variable for the appliance or for a specified queue manager
“addcert” on page 375	Add the public part of a certificate to the keystore of a specific queue manager.
“createcert” on page 376	Create a self-signed certificate for a queue manager.
“createcertrequest” on page 378	Create a certificate request for a queue manager.
“deletecert” on page 379	Delete a certificate from the keystore of a specific queue manager.
“deletecertrequest” on page 380	Delete a certificate request that was previously issued from a specific queue manager.
“detailcert” on page 381	Show detailed information about a certificate for a specific queue manager.
“detailcertrequest” on page 382	Show detailed information about a certificate request for a specific queue manager.
“keybackup” on page 383	Back up the queue manager key repository to a file.
“keyrestore” on page 384	Restore a key repository
“listcert” on page 385	List the certificates that are held in the keystore of a specific queue manager.

Table 4. Appliance commands (continued)

Command	Description
"listcertrequest" on page 386	List the certificate requests that are outstanding in the keystore of a specific queue manager.
"receivecert" on page 387	Receive a certificate signed by a Certificate Authority (CA) as the result of a previous request.
"recreatecertrequest" on page 388	Re-create a certificate request for a specific queue manager.
"usercreate" on page 370	Create user IDs for messaging users.
"userdelete" on page 371	Delete a messaging user.
"usermodify" on page 371	Modify user IDs for messaging users.
"userlist" on page 372	List the messaging users.
"groupcreate" on page 372	Add user groups for messaging users.
"groupdelete" on page 373	Delete user groups for messaging users.
"grouplist" on page 373	List groups for messaging users.
"userbackup" on page 373	Back up messaging users.
"userrestore" on page 374	Restore messaging users.

Differences between queue managers that are running on the IBM MQ Appliance and an IBM MQ installation

The IBM MQ Appliance hosts queue managers.

IBM MQ Appliance queue managers are similar in their capabilities to IBM MQ queue managers hosted on supported UNIX and Linux platforms, although some differences do exist.

Exits and services on the IBM MQ Appliance

You cannot run user code on the IBM MQ Appliance. Any attempts to create administrative objects that reference user code are rejected.

The following types of exits and services are not supported:

Channel exits
Any attempt to define or alter a channel to use an exit is rejected.

Channel auto-definition exits
Any attempt to alter a queue manager to use an exit is rejected.

Cluster workload exits
Any attempt to alter a queue manager to use an exit is rejected.

Data conversion exits
You cannot upload a data conversion exit to the appliance.

Services
Any attempt to define a service is rejected.

API exits, publish exits, and user authorization services
Stanzas about API exits, publish exits, and user authorization services cannot be added to the `qm.ini` file.

MQTT services
MQTT services cannot be used on the appliance.

JAAS security exits
JAAS security exits cannot be used with the IBM MQ Light API (see IBM MQ Light API).

The following type of service is supported:

IBM MQ Light API
From IBM MQ V8.0.0.5, the IBM MQ Light API can be used on the appliance.

Queue manager configuration on the IBM MQ Appliance

Queue managers on the IBM MQ Appliance are created with different default values to queue managers in IBM MQ.

Maximum channels

On the IBM MQ Appliance, you do not need to alter the `MaxChannels` or `MaxActiveChannels` attributes to define the maximum number of channels that can concurrently connect to a queue manager. On the IBM MQ Appliance, the default value of the `MaxChannels` and `MaxActiveChannels` attributes is set to `infinite`.

If you want to limit the maximum number of client channels, use the per-channel **`MAXINST`** and **`MAXINSTC`** attributes on the SVRCONN channel definitions to define limits for each SVRCONN channel:

MAXINST
Sets the maximum number of simultaneous instances of an individual SVRCONN channel that can be started.

MAXINSTC
Sets the maximum number of simultaneous individual SVRCONN channels that can be started from a single client. In this context, connections that originate from the same remote network address are regarded as coming from the same client.

You can use the DEFINE CHANNEL command to set these attributes. For more information, see DEFINE CHANNEL in the IBM MQ documentation.

TCP Network protocol

Any channels that are configured on an IBM MQ Appliance queue manager must be of TCP protocol type. The appliance does not support any other network protocols.

User and group permissions

IBM MQ Appliance queue managers support the user-based permissions model. Creating an authority record for a user results in only that user being granted access. To grant a group of users access, an authority record is required for the group.

Circular logging

IBM MQ Appliance queue managers support only circular logging. There is no support for creating a queue manager with linear logs.

Queue manager data

When you use the **crtmqm** command to create a queue manager on the appliance, a file system is created where all queue manager data, recovery logs, and errors logs are stored. The default size of this file system is 64 GB, but you can alter the size by using the **-fs** parameter with the **crtmqm** command.

Applications connecting to a queue manager

There are differences between applications that are connecting to a queue manager that is running on an IBM MQ Appliance, and one running in an IBM MQ installation. Queue managers that are running on the appliance support only applications that connect by using TCP, over IBM MQ channels.

IBM MQ objects on the IBM MQ Appliance

Some objects on the IBM MQ Appliance have different behavior than objects on an IBM MQ installation.

Listeners

When creating listeners on the IBM MQ Appliance, you should configure the listener to start and stop with the queue manager. You can set the listener by using the `CONTROL(QMGR)` argument to the **DEFINE LISTENER** MQSC command (see DEFINE LISTENER in the IBM MQ documentation). Alternatively, you can set the **control** property in the listener widget in the IBM MQ Console (see "Working with listeners" on page 170).

Even if you leave the **control** property of a listener set to manual, note that it will automatically stop when the queue manager stops on the appliance.

Channels

The default client channel definition table (CCDT) for a queue manager is not automatically available on the IBM MQ Appliance. Use the rcrmqobj command to generate a CCDT file for a queue manager. The generated CCDT file can be downloaded from the appliance from the mqbackup:// URI (see "Creating and downloading a CCDT file" on page 197).

For more information about client channel definition tables see Client channel definition table in the IBM MQ documentation.

XA transactions

Queue managers on an IBM MQ Appliance cannot act as XA transaction managers.

Queue managers on the appliance can participate in global transactions as resource managers, but they cannot act as transaction managers and coordinate external services.

An IBM MQ client that connects to a queue manager on an appliance is not able to issue an MQBEGIN API call. The means that the client cannot start a transaction with the queue manager acting as transaction manager.

Planning a high availability system

You can create a high availability group from a pair of IBM MQ Appliances to give your IBM MQ system a degree of resilience.

See “High availability” on page 3 for an overview of the high availability solution.

When you plan a high availability implementation, consider the following points:

- Appliances:
  - A high availability group requires two IBM MQ Appliances.
  - Both appliances should be running the same level of appliance firmware. (Appliances can operate at different levels to allow time to upgrade the appliances separately, but you should avoid configuring HA queue managers, including adding or removing queue managers, during this period.)
  - You can run queue managers on both appliances (there is no concept of an 'active' and a 'standby' appliance).
  - An appliance can run high availability queue managers or disaster recovery queue managers, and queue managers that do not belong to either group. You can also run high availability queue managers that belong to disaster recovery configurations.
  - An appliance can belong to only one HA group.
  - You can create the HA group from either of the appliances.
  - The date and time settings must be sychronized between the two appliances. You can achieve this by configuring both appliances to use the same NTP server (see “Configuring the locale, date, and time” on page 114).
- Queue managers:
  - You specify that a queue manager belongs to an HA group when you create the queue manager.
  - You can create an HA queue manager on either of the two appliances in the HA group.
  - The appliance that you create the queue manager on is the preferred appliance for that queue manager. The queue manager runs on its preferred appliance so long as that appliance is available.
  - A queue manager that belongs to a high availability group can also be set to belong to a disaster recovery configuration.
- Physical configuration:
  - The appliances in the HA group synchronize by replicating queue manager data across a 10 Gb Ethernet link. Use the eth21 interface for the replication link.
  - The appliances in the HA group are also connected by a primary and a secondary interface that both use 1 Gb Ethernet links. These interfaces are used to monitor the presence of the other appliance in the group and detect a failover situation. Use the eth13 interface for the primary link, and the eth17 interface for the secondary link.
  - For the best performance of synchronization and failover, the two appliances need to be as physically close as possible, ensuring short connecting cable lengths. However, the two appliances should not be in the same rack, in case the rack fails.

– If you connect the two appliances by using a network switch, you should use a separate switch for each of the three connections.
– If you locate the two appliances in different data centers, you should be aware the limitations outlined in "Network requirements for high availability."

The following diagram shows an example HA configuration:

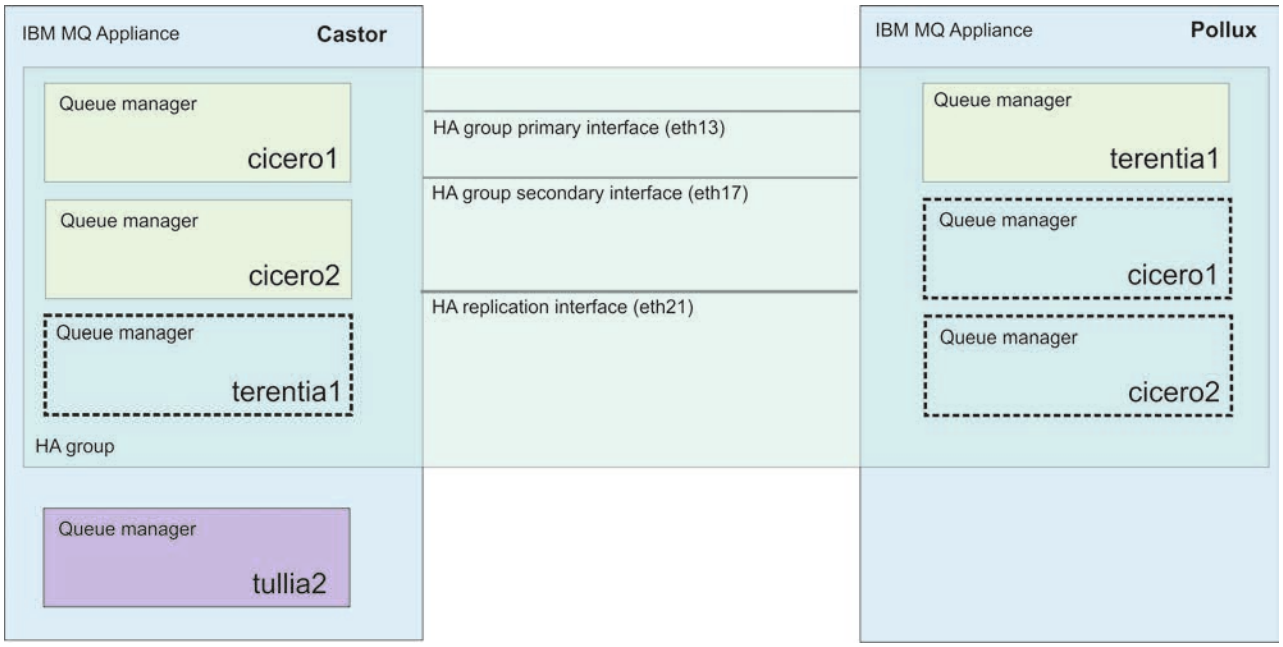

Figure 3. Example HA group

Network requirements for high availability

The high availability configuration must meet minimum network requirements to operate effectively.

Note: High availability must use the dedicated network interfaces described here for communication between the appliances (eth13, eth17 and eth21). You cannot set up a VLAN or link aggregation interface definition on the appliance in place of these named interfaces. (You can configure a VLAN on network switches, if required.)

It is recommended that you locate the appliances in a high availability pair in the same data center (preferably in adjacent racks).

If you choose to locate the appliance in different data centers, you must be aware of the following limitations:

- Performance degrades rapidly with increasing latency between data centers. Although IBM will support a latency of up to 10 ms, you might find that your application performance cannot tolerate more than 1 to 2 ms of latency.
- You must configure the primary and secondary Ethernet interfaces used for HA across completely redundant links (that is, do not rely on shared networking hardware, cabling, or power supplies for these connections). It is recommended that the replication interface is connected over a third redundant link.
- The links must have sufficient dedicated bandwidth with no contention.

- The data sent across the replication link is not subject to any additional encryption beyond that which might be in place from using MQ AMS.
- Be aware that if you lose the network connections between the two appliances, a partitioned situation can arise where the same queue manager continues to run on each appliance and each instance has a different set of queue manager data. When the connection is restored you must take action to specify which set of data you want to preserve, and which you want to discard. (This is sometimes called a 'split-brain' situation).

The design of the network topology should be performed by networking experts with a deep understanding of the network architecture being employed. The tools **ping** and **traceroute** can be used as a quick way to begin to explore the network properties, but are not a substitute for a detailed review of the network architecture.

Use **ping** to test that you can connect to the other appliance in a high availability pair:

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type the following command to test your connection:

   ```
   ping remote_IP_address
   ```

   Where *remote_IP_address* is the IP address of the eth13 interface of the other appliance in the high availability pair.
3. Repeat the **ping** test with the IP addresses of the eth17 and eth21 interfaces of the other appliance in the high availability pair.

The **ping** command sends 6 Internet Control Message Protocol (ICMP) echo-request messages to the specified host with a one second interval between each message and reports the results.

Use **traceroute** to test the connection, reporting the addresses of any hosts used to make the connection, and the latency of the connection:

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type the following command to test your connection:

   ```
   traceroute remote_IP_address
   ```

   Where *remote_IP_address* the IP address of the eth13 interface of the other appliance in the high availability pair.
3. Repeat the **traceroute** command with the IP addresses of the eth17 and eth21 interfaces of the other appliance in the high availability pair.

The **traceroute** command traces the route that packets actually take to their target host. The output shows the IP address of the hops (for example, gateway or routers) and the round trip time.

Because of the load it imposes on the network, do not use this command too often during typical operations.

Planning a disaster recovery system

You can pair a queue manager that is running on a local appliance with a queue manager that you create for the purpose on a recovery appliance in a remote location. This process provides a disaster recovery solution.

See "Disaster recovery" on page 6 for an overview of the disaster recovery solution.

When you plan a disaster recovery implementation, consider the following points:

- Appliances:
  - A disaster recovery configuration requires two IBM MQ Appliances.
  - Both appliances should be running the same level of appliance firmware. (Appliances can operate at different levels to allow time to upgrade the appliances separately, but you should avoid configuring DR queue managers during this period.)
  - You run a queue manager on the main appliance, with a back-up of that queue manager ready to run on the recovery appliance.
- Queue managers:
  - You specify that an existing queue manager is to be part of a disaster recovery configuration on the main appliance. You then run a command on the recovery appliance to create a secondary instance of that queue manager.
  - A queue manager can belong to a high availability group, and also belong to a disaster recovery configuration (see "Disaster recovery for a high availability configuration" on page 6).
  - If an event occurs that interrupts the operation of the main appliance, you can start the queue manager on the recovery appliance.
  - Messaging data is replicated asynchronously between primary and secondary queue manager. When the secondary queue manager starts, some of the messaging data might be lost (because it has not been replicated before the main appliance failed).
- Physical configuration:
  - The appliances in a disaster recovery configuration synchronize by transferring queue manager data across a 10 Gb Ethernet link.

Hardware limitations for disaster recovery

The following limitations apply to appliances configured for disaster recovery.

Note: Disaster recovery must use the dedicated network interface described here for communication between the appliances (eth20). You cannot set up a VLAN or link aggregation interface definition on the appliance in place of this named interface. (You can configure a VLAN on network switches, if required.)

Network requirements for disaster recovery

- The replication link between the two appliances must use the eth20 10 Gb Ethernet interface.
- The data sent across the replication link is not subject to any additional encryption beyond that which might be in place from using MQ AMS.
- The maximum latency for the replication link is 100 ms.

- If the IP addresses used for the two eth20 ports do not belong to the same subnet (with that subnet used only for the disaster recovery configuration) then you must set up an IP route between the eth20 ports on each appliance.
- It is recommended that the remaining 10 Gb Ethernet interfaces are used for application traffic (eth22 and eth23 on an M2001 appliance are intended for application traffic).

Planning network connections

The appliance has a number of network connections. Learn how to plan for their use.

The network connections are situated on the front of an IBM MQ Appliance. The location of the network connections is shown in the following illustration.

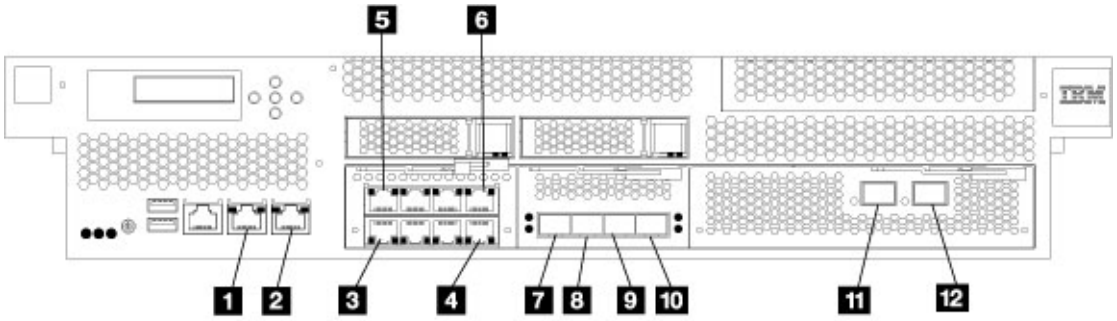

Figure 4. Network ports on an M2001 appliance

1. Management Ethernet port mgt0
2. Management Ethernet port mgt1
3. 1 Gb Ethernet module, eth10
- 1 Gb Ethernet module, eth11
- 1 Gb Ethernet module, eth12
4. 1 Gb Ethernet module, eth13
5. 1 Gb Ethernet module, eth14
- 1 Gb Ethernet module, eth15
- 1 Gb Ethernet module, eth16
6. 1 Gb Ethernet module, eth17
7. 10 Gbyte Ethernet module, eth20
8. 10 Gbyte Ethernet module, eth21
9. 10 Gbyte Ethernet module, eth22 (M2001 appliance only)
10. 10 Gbyte Ethernet module, eth23 (M2001 appliance only)
11. Fibre Channel module port 0
12. Fibre Channel module port 1

You can configure each of the interfaces when you run the installation wizard for the appliance. Alternatively, you can configure them by using the CLI or the IBM MQ Appliance web UI (see "Configuring the appliance" on page 103).

Aggregate linking and configuration for VLANs is supported for connections, except where the connections are used in a high availability or disaster recovery configuration.

IPMI LAN connection

You can use an Intelligent Platform Management Interface (IPMI) connection to interface with the Baseboard Management Controller (BMC) on the appliance. The BMC enables you to power up or power down an appliance remotely (provided that power is available in the rack where the appliance is installed). You can also use IPMI commands to perform other various functions, such as printing sensor data and reading event logs. See http://linux.die.net/man/1/ipmitool for details of IPMI commands.

You must use the mgt0 interface on the appliance for your IPMI connection. You can dedicate mgt0 to IPMI, or you can share the mgt0 connection (for example, with administration interfaces). If you share the connection, IPMI requires the allocation of a second IP address in the same VLAN and subnet as the mgt0 interface.

Because of the nature of IPMI and the level of control that can be obtained over the IBM MQ Appliance that uses this interface, either connect mgt0 to a secure management network, or do not connect it at all.

Administration

You can administer the appliance by using a command line interface (connecting via SSH) or a web user interface. You can restrict these interfaces either to mgt0 or mgt1 to create a restricted management network.

High availability

If you are implementing a high availability (HA) configuration, note that the following connections are dedicated to HA:

- eth13 - primary link
- eth17 - secondary link
- eth21 - replication link

You should have a dedicated subnet for the HA configuration, and the Ethernet connections on both appliances in the HA pair must belong to that subnet.

Disaster recovery

If you are implementing a disaster recovery (DR) configuration, note that eth20 is dedicated to replication for DR.

You should have a dedicated subnet for DR replication, and the eth20 connections on both appliances in the DR configuration must belong to that subnet.

Disaster recovery on high availability systems

If you are implementing DR on an HA system, then eth20 is used for replication by both HA appliances, and the recovery appliance.

- The IP addresses used for the two HA appliances should belong to the same dedicated subnet.
- The appliance at the recovery site does not need to belong to the same subnet, but must be able to reach it.
- The two HA appliances share a floating IP address for eth20, in addition to each having a static IP address. The floating IP is used for replication with the recovery appliance by whichever appliance is running the queue manager being replicated (that is, the primary appliance for the queue manager).
- You do not have to physically configure the floating IP. You specify it as an argument when you configure disaster recovery for an HA pair (see "Configuring disaster recovery for a high availability queue manager" on page 140). Choose an unallocated IP address on the same subnet as the two static IP addresses.

Naming interfaces

The IBM MQ Appliance enables you to specify a host alias for a specific IP address that is assigned to a network interface. You can use this alias to reference the interface, rather than explicitly using the IP address. Using aliases makes it easier to copy a configuration to another appliance, or migrate between environments, without making extensive changes to accommodate different IP addresses.

For example, the following table shows interfaces that are defined and host aliases allocated. In this example, the data connection has the host alias "data-int"; you can use this host alias in IBM MQ commands instead of explicitly referencing the IP address 10.61.121.5. The following example shows the command that is used to create a listener and bind it to the data interface:

```
define listener(CHA2L) trptype(TCP) control(QMGR) IPADDR(data-int)
```

If you then start the listener and display the status, you can see that "data-int" was resolved to IP address 10.61.121.5:

```
start listener(CHA2L)
display lsstatus(CHA2L)
AMQ8631: Display listener status details.
   LISTENER(CHA2L)                         STATUS(RUNNING)
   PID(43918)                              STARTDA(2016-05-04)
   STARTTI(09.31.56)                       DESCR( )
   TRPTYPE(TCP)                            CONTROL(QMGR)
   IPADDR(10.61.121.5)                     PORT(1414)
   BACKLOG(100)
```

For more information about defining host aliases, see "Host Alias commands" on page 530.

Table 5.

Interface	IP address	Comment	Host alias
eth10	10.61.121.5/24	Used for IBM MQ data	data-int: 10.61.121.5
eth11	10.161.121.5/25	Used for logging	log-int: 10.161.121.5
eth13	192.168.121.5/29	Used for HA Primary	hap-int: 192.168.121.5

Table 5. (continued)

Interface	IP address	Comment	Host alias
eth17	192.168.122.5/29	Used for HA Secondary	has-int: 192.168.122.5
eth21	192.168.123.5/29	Used for replication	har-int: 192.168.123.5

Data connections

You must configure one or more Ethernet connections for the IBM MQ data handled by the appliance. You can use link aggregation to improve the resilience and bandwidth of your data connection.

The M2001 appliance has two 10 Gb connections, eth22 and eth23, that are not used for HA or DR. You can aggregate these two links together to act as a single data interface.

Example

The following table shows the network configuration of an example appliance. The appliance is part of an HA group, and also supports disaster recovery for queue managers. IBM MQ data is carried on link aggregated 10 Gb connections, logging data is sent to link aggregated 1 Gb connections.

Table 6. Example network configuration for M2001 appliance

Interface	Used for
mgt0	Web UI and IPMI
mgt1	Command line access (SSH)
eth10	Not used
eth11	Not used
eth12	Link aggregated 1 Gb interface for logging data
eth13	HA primary group interface
eth14	Not used
eth15	Not used
eth16	Link aggregated 1 Gb interface for logging data
eth17	HA group alternative interface
eth20	DR replication link
eth21	HA replication interface
eth22	Link aggregated 10 Gb data interface
eth23	Link aggregated 10 Gb data interface

Capacity planning

There is no architectural limit on the IBM MQ Appliance as to the number of queue managers which can be hosted on a single appliance. However, there are some guidelines to upper limits based on the capacity of the appliance hardware and IBM tested configurations.

- For stand-alone queue managers, observe a maximum of 60 queue managers on a single appliance.

- For queue managers configured for high availability or disaster recovery (or both), observe a maximum of 30 queue manager per HA or DR pair of appliances.
- For an appliance hosting a mixture of HA or DR queue managers, and stand-alone queue managers, estimate the maximum by assuming that HA/DR queue managers consume approximately twice the resources that stand-alone queue managers consume. For example, you could plan a system with a maximum capacity of 40 stand-alone queue managers and ten HA queue managers.

These guidelines assume relatively light loading on each queue manager. For heavily loaded queue managers you might not achieve acceptable performance in practice at these levels of co-tenancy. One heavily loaded queue manager is capable of consuming all resources in the system, dependent on your configuration. For more information on this, and other performance and capacity characteristics, refer to the appliance performance report documents:

- IBM MQ Appliance Performance Report
- IBM MQ Appliance HA/DR Performance Report

Chapter 3. Installing

Plan for your installation, install, and verify the installation of the IBM MQ Appliance.

Safety

Before you install this product, read the Safety Information.

Arabic

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الآمنية

Brazilian Portuguese

Antes de instalar este produto, leia as Informações de Segurança.

Chinese (simplified)

在安装本产品之前，请仔细阅读 Safety Information
(安全信息)。

Chinese (traditional)

安裝本產品之前，請先閱讀「安全資訊」。

Croatian

Prije instalacije ovog produkta obavezno pročitajte Sigurnosne Upute.

Czech Před instalací tohoto produktu si přečtěte příručku bezpečnostních instrukcí.

Danish

Læs sikkerhedsforskrifterne, før du installerer dette produkt.

Dutch Lees voordat u dit product installeert eerst de veiligheidsvoorschriften.

Finnish

Ennen kuin asennat tämän tuotten, lue turvaohjeet kohdasta Safety Information.

French

Avant d'installer ce produit, lisez les consignes de sécurité.

German

Vor der Installation dieses Produkts die Sicherheitshinweise lesen.

Greek Πριν εγκαταστήσετε το προϊόν αυτό, διαβάστε τις πληροφορίες ασφάλειας (safety information).

Hebrew

לפני שתתקינו מוצר זה, קראו את הוראות הבטיחות.

Hungarian

A termék telepítése előtt olvassa el a Biztonsági előírásokat!

Italian Prima di installare questo prodotto, leggere le Informazioni sulla Sicurezza.

Japanese

製品の設置の前に、安全情報をお読みください。

Korean

본 제품을 설치하기 전에 안전 정보를 읽으십시오.

© Copyright IBM Corp. 2015, 2016

Macedonian

Пред да се инсталира овој продукт, прочитајте информацијата за безбедност.

Norwegian

Les sikkerhetsinformasjonen (Safety Information) før du installerer dette produktet.

Polish Przed zainstalowaniem tego produktu, należy zapoznać się z książką "Informacje dotyczące bezpieczeństwa" (Safety Information).

Portuguese

Antes de instalar este produto, leia as Informações sobre Segurança.

Russian

Перед установкой продукта прочтите инструкции по технике безопасности.

Slovak

Pred inštaláciou tohto zariadenia si pečítaje Bezpečnostné predpisy.

Slovenian

Pred namestitvijo tega proizvoda preberite Varnostne informacije.

Spanish

Antes de instalar este producto, lea la información seguridad.

Swedish

Läs säkerhetsinformationen innan du installerar den här produkten.

Guidelines for servicing electrical equipment

You must observe the guidelines when you service electrical equipment.

For your safety, the following guidelines must be observed:

- Check the area for electrical hazards, such as moist floors, non-grounded power extension cords, and missing safety grounds.
- Use only approved tools and test equipment. Some hand tools have handles that are covered with a soft material that does not provide insulation from live electrical current.
- Regularly inspect and maintain your electrical hand tools for safe operational condition. Do not use worn or broken tools or testers.
- Do not touch the reflective surface of a dental mirror to a live electrical circuit. The surface is conductive and can cause personal injury or equipment damage if it touches a live electrical circuit.
- Some rubber floor mats contain small conductive fibers to decrease electrostatic discharge. Do not use this type of mat to protect yourself from electrical shock.
- Do not work alone under hazardous conditions or near equipment that has hazardous voltages.
- Locate the emergency power-off (EPO) switch, disconnecting switch, or electrical outlet so that you can turn off the power quickly in the event of an electrical accident.
- Disconnect all power before you conduct a mechanical inspection, work near power supplies, or remove or install main units.
- Before you work on the equipment, disconnect the power cord. If you cannot disconnect the power cord, have the customer power off the wall box that supplies power to the equipment and lock the wall box in the off position.

- Never assume that power is disconnected from a circuit. Check the circuit to make sure that power is disconnected.
- If you must work on equipment with exposed electrical circuits, observe the following precautions:
  - Make sure that another person who is familiar with the power-off controls is near you and is available to turn off the power if necessary.
  - When you are working with powered-on electrical equipment, use only one hand. Keep the other hand in your pocket or behind your back to avoid creating a complete circuit that might cause an electrical shock.
  - When you use a circuit tester, set the controls correctly and use the approved probe leads and accessories for the device.
  - Stand on a suitable rubber mat to insulate you from grounds such as metal floor strips and equipment frames.
- Use extreme care when you measure high voltages.
- To ensure proper grounding of components, such as power supplies, pumps, blowers, fans, and motor generators, do not service these components outside of their normal operating locations.
- If an electrical accident occurs, use caution, turn off the power, and send another person to get medical aid.

Inspecting for unsafe conditions

How to identify potentially unsafe conditions in an IBM product that you are working on.

About this task

Each IBM product, as it was designed and manufactured, possesses safety requirements to protect users and service technicians from injury. Use good judgment to identify potentially unsafe conditions that might be caused by attachment of non-IBM features or options that are not addressed in the documentation. If you identify an unsafe condition, you must determine how serious the hazard is and whether you must correct the problem before you work on the product.

Consider the following conditions, and the safety hazards that they present:

- Electrical hazards (especially primary power). Primary voltage on the frame can cause serious or fatal electrical shock.
- Explosive hazards, such as a damaged CRT face or a bulging capacitor.
- Mechanical hazards, such as loose or missing hardware.

Procedure

1. Make sure that the power is off and the power cords are disconnected.
2. Make sure that the exterior cover is not damaged or broken, and inspect for any sharp edges.
3. Check the power cords:
   a. Make sure that the third-wire ground connector is in good condition. Use a meter to measure third-wire ground continuity for 0.1 ohm or less between the external ground pin and the frame ground.
   b. Make sure that the power cords are the correct type.
   c. Make sure that the insulation is not frayed or worn.
4. Check for pinched cables.

Safety statements

Safety statements are available on the included CD-ROM.

The *IBM Systems: Safety Notices* document is provided on the CD-ROM provided with the system.

Safety notices

These notices apply to this product.

DANGER notices warn you of conditions or procedures that can result in death or severe personal injury. **CAUTION** notices warn you of conditions or procedures that can cause personal injury that is neither lethal nor extremely hazardous. **ATTENTION** notices warn you of conditions or procedures that can cause damage to machines, equipment, or programs.

Danger notices

The following DANGER notices apply to this product.

DANGER

To prevent a possible shock from touching two surfaces with different protective ground (earth), use one hand when possible to connect or disconnect signal cables. (D001)

DANGER

Overloading a branch circuit is potentially a fire hazard and a shock hazard under certain conditions. To avoid these hazards, ensure that your system electrical requirements do not exceed branch circuit protection requirements. Refer to the information that is provided with your device or the power rating label for electrical specifications. (D002)

DANGER

If the receptacle has a metal shell, do not touch the shell until you complete the voltage and grounding checks. Improper wiring or grounding might place dangerous voltage on the metal shell. If any of the conditions are not as described, *stop*. Ensure that the proper voltage or impedance conditions are corrected before proceeding. (D003)

DANGER

An electrical outlet that is not correctly wired might place hazardous voltage on the metal parts of the system or devices that attach to the system. The customer is responsible to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)

DANGER

When you work on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- **Connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.**
- **Do not open or service any power supply assembly.**
- **Do not connect or disconnect any cables or install, maintain, or reconfigure this product during an electrical storm.**
- **The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.**
- **Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.**
- **Connect any equipment that is attached to this product to properly wired outlets.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following procedures when you install, move, or open covers on this product or attached devices.**

To disconnect:

1. **Turn off everything (unless instructed otherwise).**
2. **Remove the power cords from the outlets.**
3. **Remove the signal cables from the connectors.**
4. **Remove all cables from devices.**

To connect:

1. **Turn off everything (unless instructed otherwise).**
2. **Attach all cables to devices.**
3. **Attach the signal cables to the connectors.**
4. **Attach the power cords to the outlets.**
5. **Turn on the devices.**

(D005)

Caution notices

The following caution notices apply to this product.

CAUTION:

- **Do not install a unit in a rack where the internal rack ambient temperature exceeds what the manufacturer recommends for each of your rack-mounted devices.**
- **Do not install devices in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a component that is used for air flow through the unit.**
- **Pay attention to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels on each piece of equipment in the rack and determine the total power requirement of the supply circuit.**

- **For sliding drawers, do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.**
- **Fixed drawers must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2)**

CAUTION:
The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do not:

- **Drop or immerse into water**
- **Heat to more than 100° C (212° F)**
- **Repair or disassemble**

Exchange only with the part approved by IBM. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)

Laser safety information

CAUTION:
This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- **Do not remove the covers. Removing the covers of the laser product can result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.**
- **Use of the controls or adjustments, or performance of procedures other than what the instructions specify, can result in hazardous radiation exposure. (C026)**

CAUTION:
Data processing environments can contain equipment that transmits or receives data with laser modules that operate at greater than Class 1 power levels. To prevent permanent injury, never look into the end of an optical fiber cable or open receptacle. (C027)

Product handling information

CAUTION:

The weight of this part or unit is 18 - 32 kg (39.7 - 70.5 lb). It takes two persons to safely lift this part or unit. (C009)

Labels

One or more of the following safety labels may apply to this product.

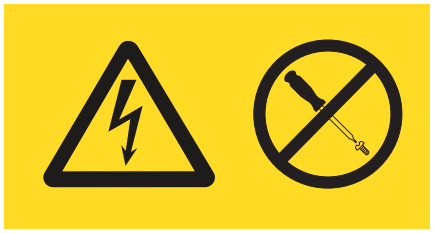

DANGER

Hazardous voltage, current, or energy levels are present inside. Do not open any cover or barrier. (L001)

DANGER

Rack-mounted devices are not to be used as shelves or work spaces. (L002)

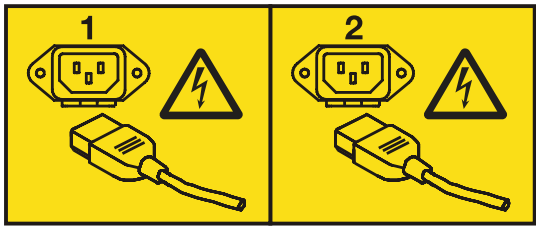

DANGER

Multiple power cords. The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. (L003)

Introducing the IBM MQ Appliance

The IBM MQ Appliance provides IBM MQ V8 on an easy-to-deploy appliance.

The IBM MQ Appliance has the model number M2001, and the MTM 8436-55X.

For completeness, details of IBM MQ Appliance M2000, MTM 8436-54X are also included here where the two models differ.

Specifications and features

This section contains information about the specifications and features of the appliances.

Specifications

Hardware specifications for the appliance.

The following table summarizes the specifications for the chassis.

Table 7. Hardware specifications

Specification	Value
Dimensions:	
Height	3.5 in. (89 mm)
Width	17.25 in. (438 mm)
Depth	23 in. (584 mm)
Appliance weight	44 lb. (20 kg)
Shipping weight	66 lb. (30 kg)
Electrical input:	
Power Supply	Two, 720 Watt power supply modules
Sine-wave	50/60 Hz (single-phase) required
110 Voltage AC	100 to 127 Volt (nominal) at 10.0 Amperes
220 Voltage AC	200 to 240 Volt (nominal) at 5.0 Amperes
Heat output	
Idle	214 watts (730 Btu/hour)
Maximum	462 watts (1575 Btu/hour)
Environment	
Shipping	-40° to 140° F (-40° to 60° C)
Power off	50° to 109.4° F (10° to 43° C)
Power on	0 to 3000 ft. (0 to 914.4 m) 50° to 95° F (10° to 35° C) 3000 ft. to 7000 ft. (914.4 m to 2133.6 m): 50° to 89.6° F (10° to 32° C)
Maximum altitude	7000 ft. (2133.6 m)
Humidity	8% to 80% (noncondensing)

Hardware features

The hardware features include processor, disk space, and memory of the appliance.

The following table describes the CPU, disk space, and memory of the appliance. Disk drive modules are serial-attached SCSI (SAS) drives.

Table 8. IBM MQ Appliance hardware features M2001

CPU	Disk space	Memory
Two 10-core 2.80 GHz Intel Xeon E5-2680V2 processors	Two 3200 GB SSDs configured as RAID 1	192 GB (Twelve 1600 MHz DDR3 DIMMs)

The machine type model (MTM) of an appliance is 8436-54X or 8436-55X.

The system disk contains 16 GB space for system file storage.

On an M2001 model the RAID array for user storage contains 3200 GB of space. Allocation of storage is set during appliance initialization.

Intrusion detection

There is an intrusion detection switch inside of the appliance.

The intrusion switch and intrusion detection are enabled by default. An administrator can configure the appliance to ignore signals from the intrusion detection switch, or reset intrusion detection.

If intrusion detection is enabled and the appliance detects an intrusion during normal operation, then the next time the appliance restarts it will enter Fail Safe mode. An administrator can reset the intrusion detection status by entering the **clear intrusion-detected** command from the CLI.

Components on the front

The following figure shows the controls, connectors, and status indicators on the front of the appliance.

M2001 model

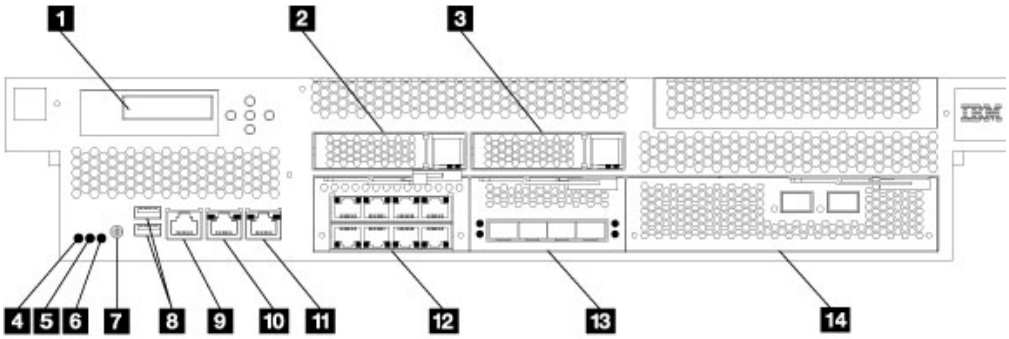

Figure 5. Controls, connectors, and status indicators on the front of the appliance.

The labels in this figure correspond to the following components on the front of the appliance:

1. LCD display.
2. Solid state disk drive 1.
3. Solid state disk drive 2.
4. Fault LED.
5. Locate LED.
6. Power LED.
7. Power button.
8. Two USB ports (not active).
9. Console connector.
10. `mgt0` management port.

11 mgt1 management port.

12 1 Gb Ethernet module.

13 10 Gb Ethernet module.

14 16 Gb fibre channel module.

M2000 model

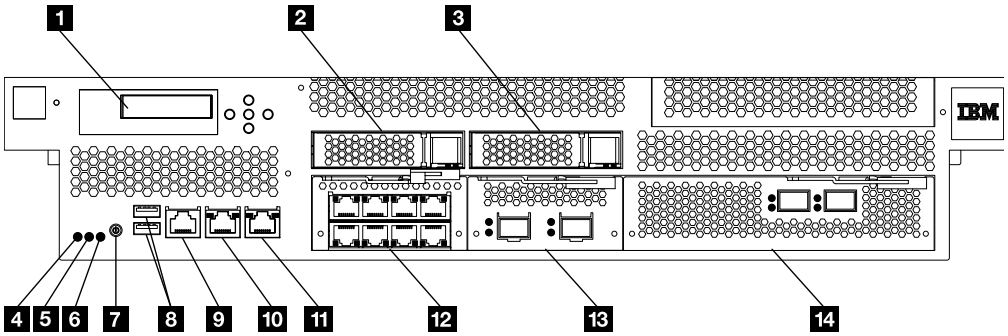

Figure 6. Controls, connectors, and status indicators on the front of the appliance.

The labels in this figure correspond to the following components on the front of the appliance:

1 LCD display.

2 Hard disk drive 1.

3 Hard disk drive 2.

4 Fault LED.

5 Locate LED.

6 Power LED.

7 Power button.

8 Two USB ports (not active).

9 Console connector.

10 mgt0 management port.

11 mgt1 management port.

12 1 Gb Ethernet module.

13 10 Gb Ethernet module.

14 16 Gb fibre channel module.

LCD module

The front panel has an LCD module that includes an LCD and five menu buttons.

The LCD displays the product name and the installed firmware version. The menu buttons adjacent to the LCD are not functional.

Locate LED

The front has a locate LED to help you identify the intended appliance.

The locate LED shows a steady blue light when activated. The LED remains on until deactivated to help you identify the intended appliance.

From the CLI

Use the **locate-device** command in Global configuration mode.

- To activate, enter `locate-device on`.
- To deactivate, enter `locate-device off`.

Power button

The front of the appliance has a power button.

When the appliance is powered off, press the button to turn on the appliance.

When the appliance is powered on, press the button to start a graceful hardware shutdown.

Console port

The front has a console port for serial communications.

The console port receives an RJ45 jack from either of the supplied serial console cables.

For initial configuration, use one of the supplied serial cables to connect from an ASCII terminal[1] to the appliance or to connect from a PC that is running terminal emulation software to the appliance.

Network ports

The network ports transmit and receive data communications between the appliance and external sources.

The network ports of a IBM MQ Appliance are grouped and located by function. Two management Ethernet ports (`mgt0` and `mgt1`) are part of the appliance. All other network ports are removable modules.

The 1 Gb Ethernet module contains eight ports for the RJ45 interface.

On the M2000 model, the 10 Gb Ethernet module has two small-form-factor pluggable (SFP+) ports.

On the M2001 model, the 10 Gb Ethernet module has four small-form-factor pluggable (SFP+) ports.

The fibre channel module has two FC ports.

1. A simple device that transmits and receives ASCII data.

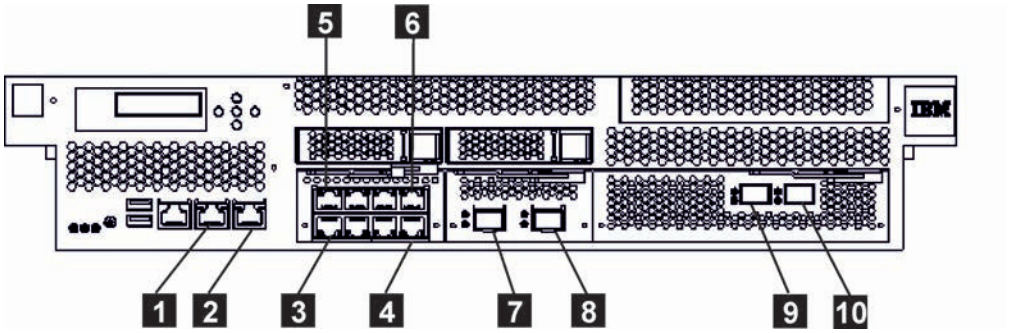

Figure 7. M2000 model network ports

1 Management Ethernet port mgt0

2 Management Ethernet port mgt1

3 1 Gb Ethernet module, eth10

- 1 Gb Ethernet module, eth11

- 1 Gb Ethernet module, eth12

4 1 Gb Ethernet module, eth13

5 1 Gb Ethernet module, eth14

- 1 Gb Ethernet module, eth15

- 1 Gb Ethernet module, eth16

6 1 Gb Ethernet module, eth17

7 10 Gbyte Ethernet module, eth20

8 10 Gbyte Ethernet module, eth21

9 Fibre channel module port 0

10 Fibre channel module port 1

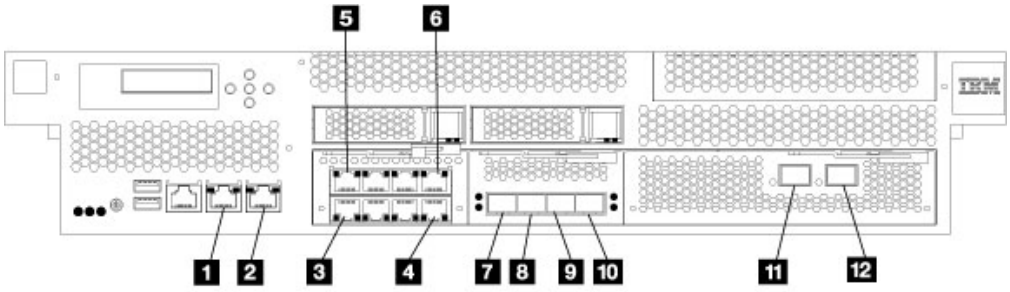

Figure 8. M2001 model network ports

1 Management Ethernet port mgt0

2 Management Ethernet port mgt1

3 1 Gb Ethernet module, eth10

- 1 Gb Ethernet module, eth11

- 1 Gb Ethernet module, eth12

4 1 Gb Ethernet module, eth13

5 1 Gb Ethernet module, eth14

- 1 Gb Ethernet module, eth15

- 1 Gb Ethernet module, eth16

6 1 Gb Ethernet module, eth17

7 10 Gbyte Ethernet module, eth20

8 10 Gbyte Ethernet module, eth21

9 10 Gbyte Ethernet module, eth22

10 10 Gbyte Ethernet module, eth23

11 Fibre channel module port 0

12 Fibre channel module port 1

Management Ethernet ports:

The `mgt0` and `mgt1` management Ethernet ports provide access to the management interfaces of the appliance.

These ports provide remote management access to the appliance and are not to be used as data ports. `mgt0` supports IPMI over LAN (including serial over LAN).

Management traffic should be considered in the overall availability, network, and management plan for the deployment. Management traffic (with the exception of IPMI) is not fundamentally different than any other kind of traffic the appliance processes. The same techniques that separate network zones apply equally to management traffic.

Ethernet modules:

The appliance contains two Ethernet modules for network connectivity.

The left module contains eight 1 Gb Ethernet ports, and the right module contains two or four 10 Gb Ethernet ports.

1 Gb Ethernet module

The 1 Gb Ethernet module contains eight ports for the RJ45 interface. The Ethernet ports are placed in two rows and are numbered sequentially from lower left to upper right. The lower row is numbered `eth10` to `eth13` and the upper row is numbered `eth14` to `eth17`. Each port has speed and activity indicator LEDs.

Notice that the speed and activity LEDs on the lower and upper rows have opposite orientation.

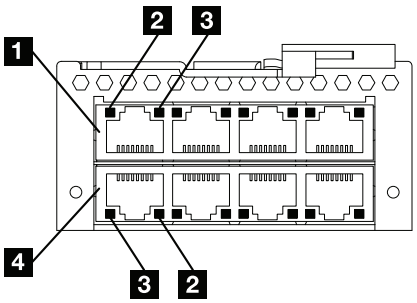

Figure 9. 1 Gb Ethernet module with eight ports for RJ45 interface

1. `eth14`
2. 1 Gb Ethernet port speed LED
3. 1 Gb Ethernet port activity LED
4. `eth10`

Use `eth13` as the HA group primary interface and `eth17` as the HA group secondary interface in a high availability pair.

4 x 10 Gb Ethernet module (M2001 models)

The 10 Gb Ethernet module has four small-form-factor pluggable (SFP+) ports. The port designators are `eth20`, `eth21`, `eth22`, and `eth23`. SFP+ ports support optical or electrical interfaces with the appropriate transceiver.

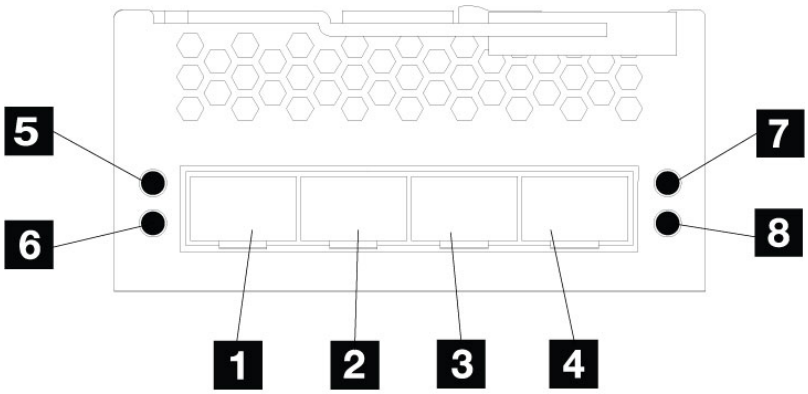

Figure 10. 10 Gb Ethernet module with two ports for SFP+ interface

1. `eth20` (used for the replication link in a disaster recovery configuration)
2. `eth21` (used for the replication link in a high availability pair)
3. `eth22`
4. `eth23`
5. `eth20` port activity LED
6. `eth21` port activity LED
7. `eth22` port activity LED
8. `eth23` port activity LED

Fibre channel module:

The appliance contains a 16 Gb fibre channel module for SAN connectivity.

The fibre channel module contains two ports, identified as port 0 and port 1. Each port has a firmware activity (green) and port activity (yellow) light.

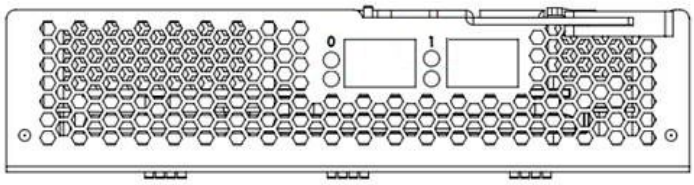

Figure 11. 16 Gb Fibre channel module

Note: The IBM MQ Appliance firmware does not currently use the fibre channel module.

Solid-state disk drive modules

The IBM MQ Appliance M2001 has two solid-state disk drive modules.

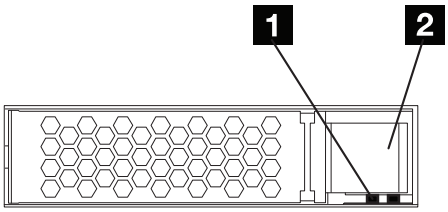

Figure 12. Solid-state disk drive module.

1 Solid-state disk drive activity LED.

2 Locking arm release latch.

Hard disk drive modules

The IBM MQ Appliance M2000 has two hard disk drive modules.

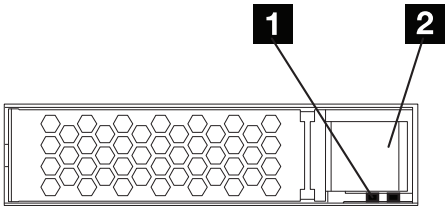

Figure 13. Hard disk drive module.

1 Hard disk drive activity LED.

2 Locking arm release latch.

Components on the rear

Fan and power supply modules are at the rear of the appliance.

The following figure shows the components at the rear of the appliance.

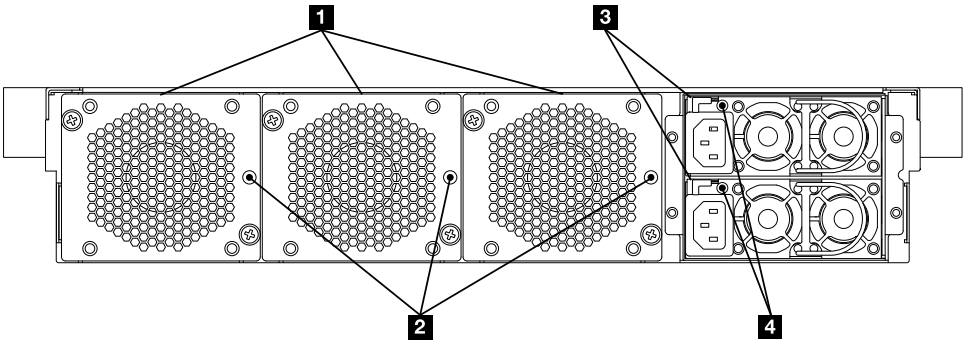

Figure 14. Rear view.

1. Fan modules.
2. Fan LEDs.
3. Power supply modules.
4. Power supply module LEDs.

The fan modules and power modules are installed from the rear of the appliance.

Fan modules

There are three fan modules in the rear of the appliance.

Each fan module contains a cooling fan with an LED that indicates the status of the module.

The speed of the fans is responsive to the temperature of the appliance as measured by internal temperature sensors near the front and rear of the appliance. As the temperature changes, the fan speed changes to compensate.

Power supply modules

The appliance is powered by two redundant power supply modules.

A single power supply module can supply the power to support appliance operations. Each power supply module contains an LED that indicates the status of the module.

DANGER

Multiple power cords. The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. (L003)

Prepare for installation

Information about the appliance, required tools, and an installation overview.

Rack requirements

Observe the rack requirements when you plan for installation.

The appliance can fit in a standard 19 in (48.26 cm) rack with a minimum of 28 in. (71.1 cm) of depth. When you plan for installation, observe the following requirements for the rack:

- The appliance rails require four mounting points in the rack.
- There must be at least 30 in. (76.20 cm) of free space behind the rack frame to remove replaceable parts.
- The ambient temperature in the operating environment and within the rack should not exceed 95° F (35° C).

DANGER

When you work on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- **Connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.**
- **Do not open or service any power supply assembly.**
- **Do not connect or disconnect any cables or install, maintain, or reconfigure this product during an electrical storm.**
- **The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.**
- **Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.**
- **Connect any equipment that is attached to this product to properly wired outlets.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following procedures when you install, move, or open covers on this product or attached devices.**

To disconnect:

1. **Turn off everything (unless instructed otherwise).**
2. **Remove the power cords from the outlets.**
3. **Remove the signal cables from the connectors.**
4. **Remove all cables from devices.**

To connect:

1. **Turn off everything (unless instructed otherwise).**
2. **Attach all cables to devices.**
3. **Attach the signal cables to the connectors.**
4. **Attach the power cords to the outlets.**
5. **Turn on the devices.**

(D005)

CAUTION:

- **Do not install a unit in a rack where the internal rack ambient temperature exceeds what the manufacturer recommends for each of your rack-mounted devices.**
- **Do not install devices in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a component that is used for air flow through the unit.**
- **Pay attention to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels on each piece of equipment in the rack and determine the total power requirement of the supply circuit.**
- **For sliding drawers, do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.**
- **Fixed drawers must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2)**

Tool requirements

You need the following tools and hardware to install the appliance rack-mounting kit.

- A medium Phillips screwdriver
- Two (2) standard rack screws

You need at least two (2) and up to 12 network cables to connect the appliance to your network.

Installing the appliance in a rack

The appliance shipping carton contains a rail kit.

The rails for the appliance are for a 19 in. (48.26 cm) rack. A complete rail kit is required to install the appliance. If any item is missing, contact IBM support.

The rail kit of the following parts:

- Left slide rail, marked L.
- Right slide rail, marked R.
- Two (2) screws (size 10-32) to secure the slide rails to the rack.

Installing rails in the rack frame

How to install rails in the rack cabinet.

Before you begin

If the rails in the kit came with thumbscrews, remove them.

Note: When you install a 2U appliance, be sure to install the slide rails in the bottom of the 2U area in the rack.

Procedure

1. Open the front rail latches, as shown in the following figure.

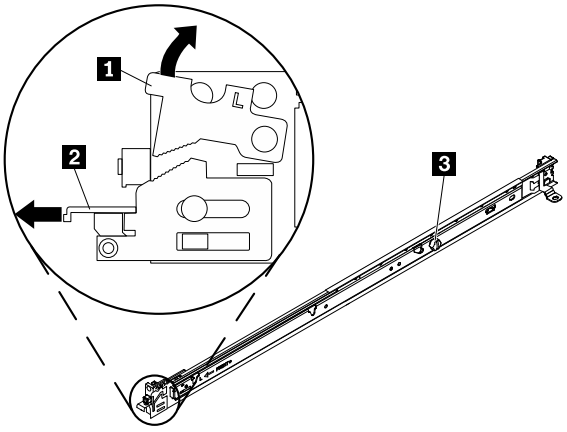

Figure 15. View of the left slide rail.

   Notice that each slide rail is marked with an R (right) or an L (left) to indicate on which side of the rack it will be installed. R and L are determined as you face the rack opening with the front portion nearest you.

   a. Select one of the slide rails and push up on the front moveable tab **1** ; then, pull out the front latch **2** .
   b. If a thumbscrew is installed in the slide rail **3** , remove it.

2. Install the rear end of the slide rails into the rack, as shown in the following figure.

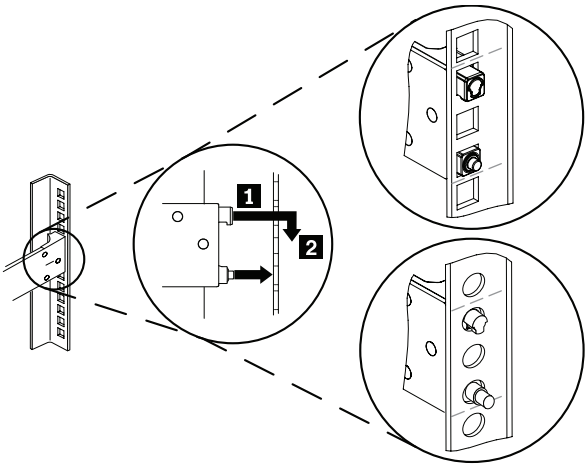

Figure 16. Install the rear end of the slide rails.

   a. From the front of the rack, line up the two pins on the rear of the slide rail with the corresponding holes at the selected location at the rear of the rack.
   b. Push the rails so that the pins go through the holes **1** , and the top pin seats into place **2** .

3. Install the front end of the rails, as shown in the following figure.

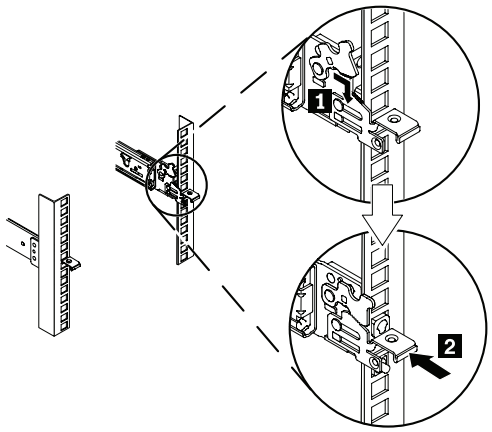

Figure 17. Install the front end of the slide rails.

a. Guide the front latch around the appropriate hole and pull the slide rail forward to fit the pins through the front of the rack.
b. Rotate the front moveable tab **1** to the downward position so that the teeth engage with the front latch.
c. Push the front latch **2** in as far as it will go.

4. Repeat steps 1 through 3 to install the other rail into the rack. Make sure that each front latch is fully engaged.
5. Install a 10-32 screw in the rear of right rail, as shown in the following figure.

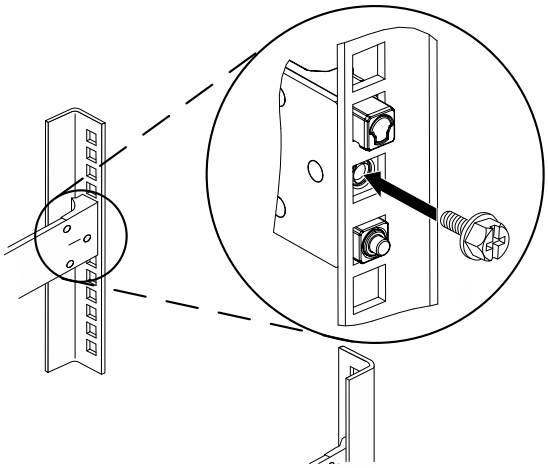

Figure 18. Securing the rails in the rack.

6. Repeat step 5 for the left rail.

Installing the appliance on the rails

How to install the appliance on the rails.

About this task

Secure the appliance to the rails. The following figure shows the numbered components that are mentioned in the steps.

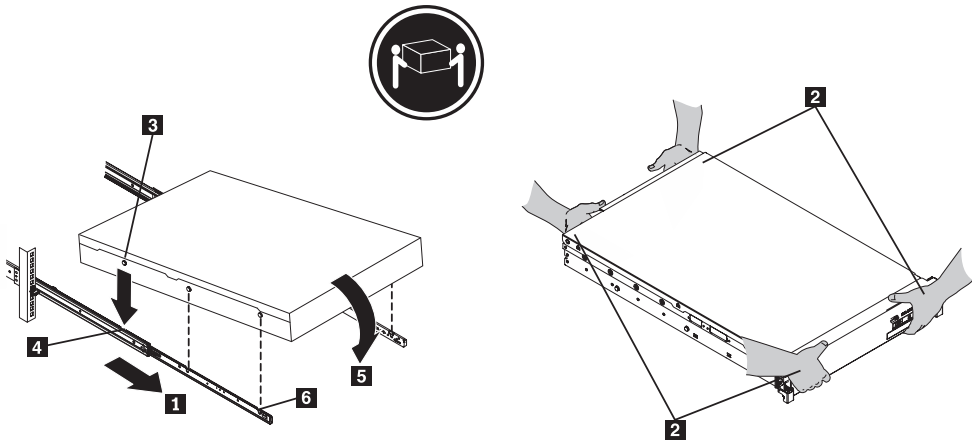

Figure 19. Securing the appliance in the rack

CAUTION:

The weight of this part or unit is 18 - 32 kg (39.7 - 70.5 lb). It takes two persons to safely lift this part or unit. (C009)

DANGER

Rack-mounted devices are not to be used as shelves or work spaces. (L002)

Procedure

1. Pull the slide rail forward **1**.
2. Use two people to carefully lift the appliance from the lifting points **2** and tilt it into position over the slide rails so that the rear nail heads **3** on the appliance line up with the rear slots **4** on the slide rails.
3. Slide the appliance down until the rear nail heads slip into the two rear slots, and then slowly lower the front of the appliance **5** until the other nail heads slip into the other slots on the slide rails.
4. Make sure that the front latch **6** slides over the nail heads.
5. Next, slide the appliance into the rack.

Sliding the appliance into the rack

Before you begin

If the appliance is locked into place, slide the appliance toward you.

About this task

The following figure shows the numbered components that are mentioned in the steps.

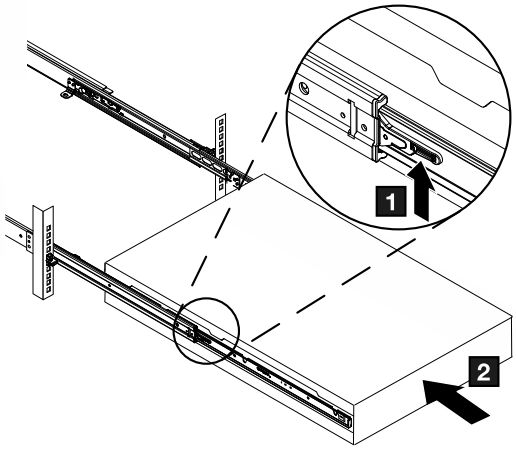

Figure 20. Sliding the appliance into the rack.

Procedure

1. Secure the brackets to the appliance with the captive screws **1** .
2. Slide the appliance into the rack **2** .

Considerations to connect the appliance to an AC power source

Read the considerations before you connect the appliance to an AC power source.

DANGER

Overloading a branch circuit is potentially a fire hazard and a shock hazard under certain conditions. To avoid these hazards, ensure that your system electrical requirements do not exceed branch circuit protection requirements. Refer to the information that is provided with your device or the power rating label for electrical specifications. (D002)

DANGER

If the receptacle has a metal shell, do not touch the shell until you complete the voltage and grounding checks. Improper wiring or grounding might place dangerous voltage on the metal shell. If any of the conditions are not as described, *stop*. Ensure that the proper voltage or impedance conditions are corrected before proceeding. (D003)

DANGER

An electrical outlet that is not correctly wired might place hazardous voltage on the metal parts of the system or devices that attach to the system. The customer is responsible to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)

You must use the provided power cords to connect both power supply modules to an AC power source. An unconnected module is considered by the system to be in a failed state.

Connect the appliance to a network

Considerations for when you connect the appliance to a network.

DANGER

To prevent a possible shock from touching two surfaces with different protective ground (earth), use one hand when possible to connect or disconnect signal cables. (D001)

CAUTION:
Data processing environments can contain equipment that transmits or receives data with laser modules that operate at greater than Class 1 power levels. To prevent permanent injury, never look into the end of an optical fiber cable or open receptacle. (C027)

Attention: Never connect the appliance to telephone or other telecommunication circuits.

Do not use a fiber optic cable that is longer than 100 meters. The cables for small-form factor pluggable (SFP+) modules can be longer than 100 meters. See the product documentation for detailed information on SFP+ modules.

The appliance Ethernet ports must be connected to a compatible link partner, preferably set to auto-negotiate connection speed and mode (half duplex or full duplex). Depending on the negotiated or static connection speed and mode, ensure that the cable complies with the following requirements:

10BASE-T (10 Mbps) connection
Two pairs of Category 3 wiring or better.

100BASE-TX (100 Mbps) connection
Two pairs of Category 5 wiring or better.

1000BASE-T (1 GbE) connection
Four pairs of Category 5 wiring or better.

10GBASE (10 Gbps) connection:

- **Short-reach (300 meters) SFP+ modules with LC connector (multi-mode fiber)**
  - Optical interface specifications per IEEE 802.3ae 10GBASE-SR
  - Mechanical specifications per SFF Committee SFF 8432 Improved Pluggable Formfactor IPF
  - Class 1 Eye safe per requirements of IEC 60825-1 / CDRH
- **Long-reach (10 km) SFP+ modules with LC connector (single-mode fiber)**
  - Optical interface specifications per IEEE 802.3ae 10GBASE-LR
  - LC Duplex optical connector interface confirming to ANSI TIA/EA 604-10 (FOCIS 10A)
  - Class 1 Eye safe per requirements of IEC 60825-1 /CDRH
- **SFP+ Copper Direct Attach twinaxial cables**

The appliance is provided with the following cables, which can be used to connect the two appliances in a high availability pair:

- Two Cat5e cables that can be used to connect ports `eth13` and `eth17` on one appliance to ports `eth13` and `eth17` on the other appliance in an HA pair.

- One Direct Attach Copper cable with SFP+ connectors that can be used to connect port `eth21` on one appliance to port `eth21` on the other appliance in an HA pair.

Setting up the initial firmware configuration

How to perform the initial, base firmware configuration.

About this task

This configuration is the minimal configuration to add an appliance to your environment. Defining the full configuration for your appliance is described elsewhere.

Procedure

1. Read the hardware and information requirements, and read the considerations for the operation modes and the password for the `admin` account. (see "Considerations for the password of the `admin` account" on page 55 and "Appliance modes" on page 55)
2. Connect the serial cable to the appliance.
3. Initialize the appliance by changing the password for the `admin` account and interactively defining the base configuration.
4. Accept the license agreement and verify the base configuration.

Configuration requirements

You must meet both hardware and information requirements to perform the initial firmware configuration.

Before you begin the initial firmware configuration, make sure that you meet the following requirements:

- You review and comply with the hardware requirements.
- You obtain the required network data.

Hardware requirement

You must use a serial connection to perform the initial configuration.

The package contains a USB serial console cable (USB to RJ45) and a DE-9 serial console cable (DE-9 to RJ45). For initial configuration, use a supplied cable to connect from an ASCII terminal to the appliance or to connect from a PC that is running terminal emulation software to the appliance.

Information requirements

Before you define the base configuration, obtain the essential network data from your network administrator.

You need IP address information for each of the following:

- Ethernet interfaces that are used for appliance management ports `mgt0` and `mgt1`.
- Ethernet interfaces that are used for service access.
- Default gateways (routers) that support the subnets for the Ethernet interfaces.
- The IP addresses and ports for the web management interface and SSH service.

Tip:

- The IBM MQ Appliance web UI is required to accept the license agreement.

- If you want to use an IPMI connection (including serial over LAN), it must be configured on `mgt0`.

Firmware considerations

During the initial firmware configuration, the script prompts you for supported operational modes and the password for the `admin` account.

Considerations for the password of the `admin` account

On the first boot of the appliance, you must change the password for `admin` account.

- On the first boot, you must initialize the appliance. The initialization routine prompts you to change the password for the `admin` account. Then, you are prompted to create a user of the privileged account type, or the group-defined account type (with the appropriate access policy) as a backup for the `admin` account. A privileged, or group-defined user (with the appropriate access policy) can log in and reset the password for the `admin` account.
- On subsequent boots, you are prompted for the credentials of the `admin` account or another local account. If the account password is expired, you are prompted to change the password.

Attention: Do not forget or misplace the password for the `admin` account. If you forget or misplace this password, security best practice recommends that you return the appliance to IBM to reset this password. However, if another user account can log in and has the appropriate access permission, that user can reset the password for the `admin` account. You can define additional administrative accounts, see Configuring appliance users in IBM Knowledge Center.

When you receive the appliance after a password-reset, you must perform an initial firmware setup that removes all existing configuration data from the appliance.

Appliance modes

You must confirm which mode your IBM MQ Appliance appliance operates in.

Depending on the product license purchased, the IBM MQ Appliance can operate in one of two modes:

- IBM MQ Appliance M2001A (or) is aimed at larger enterprise workloads.
- IBM MQ Appliance M2001B (or) is designed to meet the needs for smaller workloads and offers a lower processing capability.

You can verify in IBM Passport advantage which product license you have purchased for your appliance.

The first time that you power on the IBM MQ Appliance, you are asked to confirm which license you have purchased. The appropriate appliance is then applied. Please take care when making this choice. If you accidentally configure this setting incorrectly, you must apply a factory reset to the appliance (see Factory reset) or contact IBM Support.

Once configured you will not be asked again as the appliance mode has now been selected. The mode is indicated in the LCD panel on the front of the appliance and the welcome banner when logging in to the IBM MQ Appliance CLI.

If you later require more capacity, you can purchase an upgrade to convert an M2001B (or) appliance to an M2001B+ (or) appliance, which has the same capacity as an M2001A (or) appliance.

Upgrading from an M2001B to M2001B+ (or) does not require a factory reset or loss of queue manager data, but does require a reboot of the appliance for the update to take effect. Full instructions for this process are supplied on purchase of the upgrade part through passport advantage.

Procedure 1 of 3: Connecting the serial cable to the appliance

How to make the serial connection to the appliance.

Before you begin

Read the hardware and information requirements in "Configuration requirements" on page 54, and read the operation modes and password considerations for the `admin` account in "Firmware considerations" on page 55.

About this task

For initial configuration, you must connect to the appliance console port from an ASCII terminal, or a computer that is running terminal emulation software.

The DE-9 (sometimes called DB-9) serial console cable connects a 9-pin socket to an 8-position modular plug (RJ45). The cable conforms to the EIA/TIA-574 standard as data circuit-terminating equipment (DCE).

If your PC does not recognize the USB serial console cable, you might need to install a device driver. Standard drivers with installation instructions are on the *Resource Kit* in an archive file.

- The driver for Windows systems is in the `driver/win/` directory.
- The drivers for Mac OS systems are in the `driver/mac/` directory.

Notes:

- Do not connect an Ethernet network cable to the appliance serial console port.
- Do not connect a digital or analog Telephone network cable to the appliance serial console port.

DANGER

To prevent a possible shock from touching two surfaces with different protective ground (earth), use one hand when possible to connect or disconnect signal cables. (D001)

Procedure

1. Use the appropriate cable to connect from an ASCII terminal or PC that is running terminal emulation software to the appliance.
2. Ensure that the terminal or PC software is configured for standard, 115200, 8N1[2], and no flow control data transfer.

2. 8N1 is a notation for a serial configuration in asynchronous mode, where there are eight data bits, no (N) parity bit, and one stop bit.

What to do next

See "Procedure 2 of 3: Initializing the appliance" to define the base configuration such as changing the password for the `admin` account, creating privileged users, and configuring the web management interface.

Procedure 2 of 3: Initializing the appliance

Use this procedure to provide base configuration for the appliance.

Before you begin

See "Procedure 1 of 3: Connecting the serial cable to the appliance" on page 56 to connect the appliance to an ASCII terminal or to a PC that is running terminal emulation software through a serial connection.

Procedure

1. Press the power button at the front of the appliance. The green power LED illuminates.
   - You might hear the fans start.
   - You might hear the fans change speed as the screen displays `DPOS boot - press <ESC> within 7 seconds for boot options...`

   Wait for the appliance to boot.
2. At the `Login:` prompt, enter `admin`[3].
3. At the `Password:` prompt, enter `admin`[4]. The script prompts you later to change this password.
4. Follow the prompts to enable the appropriate appliance mode. Select M2001A or M2001B (or) according to your license (see "Appliance modes" on page 55).

   Attention: Use care when you select the operational modes. If you select an incorrect mode, the only way to change an operational mode is to reinitialize the appliance, which deletes all configuration settings on the appliance.
5. At the `Please enter new password:` prompt, enter a new password.
   - Ensure that your keyboard does not have Caps Lock or Number Lock engaged.
   - Type the password from the keyboard. Do not copy and paste the password. If you copy and paste, you might copy extra spaces or characters.
6. At the `Please re-enter new password to confirm:` prompt, enter the new password again.
7. At the `Do you want to run the Installation Wizard?` prompt, enter `y` to start the installation wizard.

   Note: If you inadvertently enter `n` at the prompt, you can start the installation wizard by entering the following commands:

   ```
   configure terminal
   startup
   ```
8. Follow the prompts to complete the base firmware configuration. You should configure the following features at the minimum:
   - At least one network interface for remote management.
   - The SSH service.

---

3. `admin` is the name of a local user account. The owner of this account can perform all tasks on the appliance.

4. `admin` is the default password for the `admin` account.

- The web management service. If this is not configured, you cannot accept the license agreement, and will be able to take no further actions on the appliance.
- The name for the system. This is mandatory if you are configuring the appliance as one of a high availability pair, or part of a disaster recovery configuration.

After you define the base firmware configuration, the screen displays information that is similar to the following example. The screen shows information specific to your appliance.

```
Welcome to IBM MQ appliance M2001A console configuration.
Copyright IBM Corporation 1999-2015

Version: MQ00.8.0.0.5 build 000000 on 2016/08/18 12:24:18
Serial number: DPTP004

You must read and agree to the terms of the license agreement using the WebGUI.
If you did not configure the Web Management Interface, you must do it now with
the following command:
 configure terminal;web-mgmt;admin-state enabled;local-address 0 9090;exit
mqa#
```

The previous sample shows the following information:

- The appliance is an IBM MQ Appliance.
- The firmware version that is running on the appliance is `8.0.0.5` at the `000000` build level.
- The date and time that build 000000 was created is August 18, 2016 at 12:24:18.
- The serial number of this appliance is `DPTP004`.
- Instructions to access the license agreement.

What to do next

See "Procedure 3 of 3: Accepting the license agreement" to access the IBM MQ Appliance web UI and accept the license agreement.

Procedure 3 of 3: Accepting the license agreement

You must access the IBM MQ Appliance web UI and accept the license agreement.

Before you begin

See "Procedure 2 of 3: Initializing the appliance" on page 57 to define the base configuration for the appliance.

About this task

This procedure makes the following assumptions:

- The IP address for the Ethernet interface that is used to access the IBM MQ Appliance web UI is `10.10.13.35`
- The specialized HTTP server to support IBM MQ Appliance web UI access listens on port `9090`

You can discover the IP address by using the command `show ipaddress` on the command line.

Procedure

1. Open a web browser.
2. In the **Address** field, enter `https://10.10.13.35:9090`. If the web page is displayed successfully, the base firmware configuration is successful.
3. Log in to the appliance with the local administrator account and password.
4. Click **Login**. The IBM MQ Appliance web UI displays the license agreement.
   - Click **I agree** to accept the terms of the license agreement and non-IBM terms. The appliance reloads the firmware. In a few minutes, you can log in again after the appliance restarts.
   - If you do not agree, click **I do not agree**. The initialization of the appliance stops. You need to either power off the appliance or review and accept the license agreement.
5. Log in again to verify that the `admin` account and other administrators can access the appliance with their credentials.

What to do next

To access the information about completing the configuration beyond the base configuration, such as creating additional users and groups, configuring interfaces, setting up high availability and so on, see Configuring in IBM Knowledge Center.

Maintenance

Topics in this section describe general fault finding and maintenance of the appliance.

Diagnosing your appliance

How to diagnose problems in your appliance.

Before you perform maintenance on this product, read the safety information.

Use the indication of LEDs, **`test hardware`** command, diagnostic self-test, and status providers for the sensors to diagnose problems with the appliance and modules.

Appliance LEDs

LEDs help you diagnose possible problems with the hardware components of an appliance.

You can use the following LEDs to determine the behavior and diagnose a problem with the appliance and components:

- Fault LED, locate LED, and power LED at the front of the appliance.
- Activity and speed LEDs of Ethernet modules.
- Activity LEDs of hard disk drive modules.
- LEDs of fan modules.
- LEDs of power supply modules.

LEDs on the front of the appliance:

The following figure describes LEDs of the appliance.

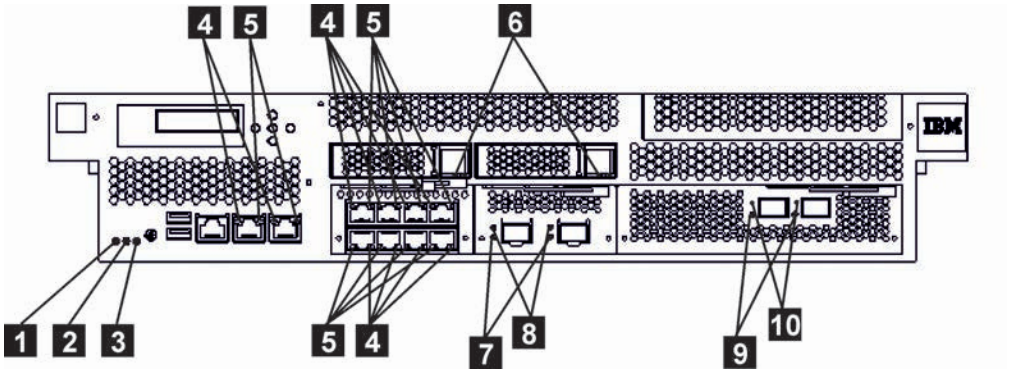

Figure 21. LEDs on the front of the M2000 appliance

The labels in this figure correspond to the following LEDs on the front of the appliance:

1 Fault LED.

This indicator shows steady amber light when the appliance detects a critical hardware event.

2 Locate LED.

This indicator shows steady blue light when activated.

3 Power LED.

This indicator shows green steady light when the power is connected and the appliance is turned on.

4 1 Gb Ethernet port speed LED

Green steady light indicates a 1 Gb Ethernet connection.

Amber steady light indicates a 10 or 100 Mbps connection.

5 1 Gb Ethernet port activity LED

Green steady light indicates when the port is connected.

Green flashing light corresponds to port activity.

6 Hard disk drive activity LED

Green steady light is present when the module is inserted fully.

Green flashing light corresponds to the reading or writing of data on the disk.

7 10 Gb Ethernet port speed LED

Green steady light indicates a 1 Gb Ethernet connection.

Amber steady light indicates a 10 Gb Ethernet connection.

8 10 Gb Ethernet port activity LED

Green steady light indicates when the Ethernet port is connected.

Green flashing light corresponds to port activity.

9 16 Gb fibre channel port activity LED

Yellow - two fast flashes indicates a 4 Gbps link rate.

Yellow - three fast flashes indicates a 8 Gbps link rate.

Yellow - four fast flashes indicates a 16 Gbps link rate.

10 16 Gb fibre channel firmware activity LED

Green steady light indicates when the link is active.

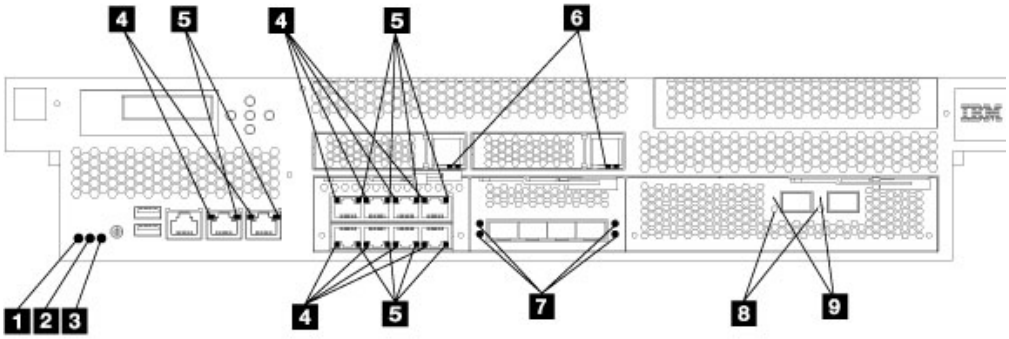

Figure 22. LEDs on the front of the M2001 appliance

The labels in this figure correspond to the following LEDs on the front of the appliance:

1 Fault LED.

This indicator shows steady amber light when the appliance detects a critical hardware event.

2 Locate LED.

This indicator shows steady blue light when activated.

3 Power LED.

This indicator shows green steady light when the power is connected and the appliance is turned on.

4 1 Gb Ethernet port speed LED

Green steady light indicates a 1 Gb Ethernet connection.

Amber steady light indicates a 10 or 100 Mbps connection.

5 1 Gb Ethernet port activity LED

Green steady light indicates when the port is connected.

Green flashing light corresponds to port activity.

6 Solid state disk drive activity LED

Green steady light is present when the module is inserted fully.

Green flashing light corresponds to the reading or writing of data on the disk.

7 10 Gb Ethernet port speed LED

Green steady light indicates a 1 Gb Ethernet connection.

Amber steady light indicates a 10 Gb Ethernet connection.

8 16 Gb fibre channel port activity LED

Yellow - two fast flashes indicates a 4 Gbps link rate.

Yellow - three fast flashes indicates a 8 Gbps link rate.

Yellow - four fast flashes indicates a 16 Gbps link rate.

9 16 Gb fibre channel firmware activity LED

Green steady light indicates when the link is active.

LEDs on the rear of the appliance:

The LEDs on the rear panel of the appliance provide diagnostic information about power supply and fan modules.

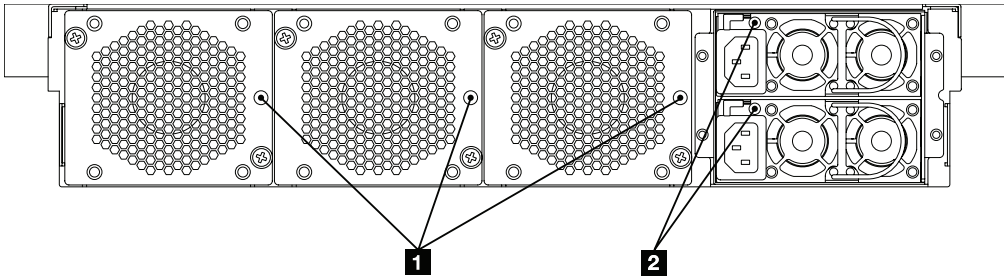

Figure 23. LEDs on the rear of the appliance

The labels in this figure correspond to the following LEDs on the rear of the appliance:

1 Fan LEDs.

- Amber single flash shows when power is first applied to the fan module.
- Amber steady light indicates that the fan is operating at less than 1200 revolutions per minute (RPM) or there is a fault in the module.
- No illumination when there is no power present or there is no problem.

2 Power module LEDs.

- Green steady light indicates that the module is connected to a power source.
- Red steady light indicates that the module is not functioning within design specifications.
- If not illuminated, there is no power to the module.

`test hardware` command

You can use the Global **`test hardware`** command to test the hardware from the CLI.

To use this command:

- You must establish an SSH connection to the appliance.
- You are in Global configuration mode (set with the **`configure terminal`** command).

To test the hardware from the configuration, enter the following commands:

```
# configure terminal
(config)# test hardware
```

Depending on the state of the hardware, the command produces output that shows the status of each component:

- `success`
- `warning`
- `failure`

The components are broken down into the following categories:

- Backtrace availability
- Interface diagnostics
- Fan diagnostics
- Cryptographic card diagnostics
- RAID volume diagnostics
- Sensors diagnostics
- CPU/memory diagnostics

Samples of `success` statements are as follows:

- ```
  [success] Status of voltage reading 'Voltage PU +12' : ok.
  [success] Status of voltage reading 'Voltage PU +3.3' : ok.
  [success] Status of voltage reading 'Voltage PU +5' : ok.
  ```
- ```
  [success] CPUs OK
  [success] Memory all present
   DIMM_A1 0x0015 16384 MB Micron             36KSF2G72PZ-1 0C676D47
   DIMM_A2 0x0015 16384 MB Micron             36KSF2G72PZ-1 0C676D62
   DIMM_B1 0x0015 16384 MB Micron             36KSF2G72PZ-1 0C676C08
   DIMM_B2 0x0015 16384 MB Micron             36KSF2G72PZ-1 0C676B80
   DIMM_C1 0x0015 16384 MB Micron             36KSF2G72PZ-1 0C676C91
   DIMM_C2 0x0015 16384 MB Micron             36KSF2G72PZ-1 0C676C59
   DIMM_D1 0x0015 16384 MB Micron             36KSF2G72PZ-1 0C676BCD
   DIMM_D2 0x0015 16384 MB Micron             36KSF2G72PZ-1 0C676C71
   DIMM_E1 0x001F 16384 MB Micron             36KSF2G72PZ-1 0C676D68
   DIMM_F1 0x001F 16384 MB Micron             36KSF2G72PZ-1 0C676B99
   DIMM_G1 0x001F 16384 MB Micron             36KSF2G72PZ-1 0C676C68
   DIMM_H1 0x001F 16384 MB Micron             36KSF2G72PZ-1 0C676CE1
  ```
- `[success] Statistics for interface 'eth10' show no errors`
- `[success] fan 1 operating within expected range`
- `[success] Status of crypto 'hardware2' : fully operational.`

Samples of `warning` statements are as follows:

- `[warning] No RAID Battery Backup Unit found.`
- `[warning] Physical link on interface 'eth10' is down.`
- `[warning] eth10 has invalid MAC (ff:ff:ff:ff:ff)`

Samples of `failure` statements are as follows:

- `[failure] Memory in error DIMM_H1, 0x001F`
- `[failure] fan 2 operating outside expected range (rpm too low)`
- `[failure] Status of crypto 'not detected' is unknown.`

The output of the **`test hardware`** command is part of any generated error report.

Using the diagnostic self-test

The appliance provides a boot-time diagnostic self-test to help you test hardware components.

About this task

Only use the diagnostic self-test when directed by IBM Support to help confirm a potential hardware problem with the appliance.

Procedure

1. Connect the serial cable.
2. If the appliance is not turned on, press the power button to turn on the appliance. The green power LED illuminates. You should hear the fans start.
3. When you see `DPOS boot - press <ESC> within 7 seconds for boot options`, press ESC. You should see the `DPOS` prompt followed by the boot options menu.

```
DPOS boot - press <ESC> within 7 seconds for boot options.. <ESC>
DPOS> ?
Available boot options:

Boot Option   Description
------------  ----------------------------------------
system        Normal System Startup
diagnostics   Run Standalone Hardware Diagnostics

DPOS>
```

4. At the `DPOS` prompt, enter `diagnostics` to start the appliance and display the diagnostics main menu.

```
Hardware Diagnostics Tool Version 1.0
(C) Copyright 2011, 2014 - IBM Corporation

Main Menu:
    1. Inventory                                        n/a
    2. BMC/Sensors                                      n/a
    3. Network                                          n/a
    4. Memory                                           n/a
    5. Disks                                            n/a
    0. Exit Diagnostics

Select action>
```

5. To select a test to run, enter its number at the `Select action` prompt.

Results

After a test completes, the diagnostic self-test produces one of the following results:

- `PASS`
- `FAIL`
- `RUNNING`
- `SKIP`
- `n/a`

Viewing status providers for sensors

This section introduces the status providers for sensors that monitor the components of the appliance.

The appliance provides the following sensors status providers:

Fan speed sensors
Provides the measured speed in RPM for the fans in each fan module. You can view the results of the fan speed sensors from the the CLI, enter **show sensors-fans**.

Temperature sensors
Provides the measured temperature in degrees Celsius for internal components:

- Temperature of each CPU and each DIMM of the CPU components
- Air temperature
  - The System 1 sensor reads the temperature at the front of the appliance.
  - The System 2 sensor reads the temperature at the rear of the appliance.

You can view the results of the temperature sensors from the CLI, enter **show sensors-temperature**. The temperature is in degree Celsius.

Voltage sensors
Provides the measured voltage for the components in millivolts. You can view the results of the voltage sensors from the CLI, enter **show sensors-voltage**.:

Current sensors
Provides the measured current for the internal components in milliamperes. You can view the results of the current sensors from the CLI, enter **show sensors-current**.

RAID battery backup status
Monitors the power backup unit connected to the RAID controller. You can view the RAID battery backup status from the CLI, enter **show raid-battery-module**.

Other sensors
Provides Boolean values for the status of intrusion switch and power supply modules.

- A value of `true` indicates that the condition exists.
- A value of `false` indicates that the conditions does not exist.
- For the intrusion switch, the value indicates whether it was tripped.
- For each power supply, the value indicates the condition:
  - Output Failure: The power supply module failed.
  - AC lost: The power cord is not attached.
- For each hard disk in the array and the battery, the values indicates the state:
  - Fault
  - Present

You can view the results of the other sensors from the CLI, enter **show sensors-other**.

Troubleshooting your appliance

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and to explain how to resolve the problem.

Follow the troubleshooting workflow to troubleshoot hardware problems with the appliance.

Troubleshooting workflow

Use this workflow to troubleshoot the problem and determine whether you need to contact IBM Support for assistance or to order a replacement part.

Procedure

1. Did you receive a critical event through SMTP notification?

   The following messages are examples of critical messages:

   - `[system][critic] sensors: tid(id): System power supply number has failed.`
   - `[system][critic] sensors-fans: tid(id): Chassis cooling fan number operating too slowly.`

   For information about creating log targets for notification, see the managing logs topic.

   Yes Continue to step 3.

   No Continue to step 2.

2. Does the log file contain a critical message?

   For information about viewing logs, see the viewing logs topic.

   Yes Continue to step 3.

   No Continue to step 4.

3. Does the critical event or critical log message identify the part that is failing or has failed?

   Yes Continue troubleshooting to determine whether you need a replacement part:

   - If a fan module, see "Troubleshooting fan modules."
   - If the power supply module, see "Troubleshooting power supply modules" on page 67
   - If the hard disk drive module, see "Troubleshooting disk drive modules" on page 67.
   - If field replaceable unit (FRU) parts, contact IBM Support.

   No Continue to step 4.

4. Is the Fault LED illuminated on the front of the appliance?

   Yes Continue with step 5.

   No The problem is with the appliance, use the appliance troubleshooting procedure.

5. Are the LEDs lit for any modules?

   Yes

   If a fan module, see "Troubleshooting fan modules."

   If the power supply module, see "Troubleshooting power supply modules" on page 67

   If the hard disk drive module, see "Troubleshooting disk drive modules" on page 67.

   No The problem is with the appliance, use the appliance troubleshooting procedure.

Troubleshooting fan modules

How to troubleshoot the fan modules.

About this task

When one or more fans are not working, turn off the appliance as soon as possible to avoid overheating. The remaining fans might not be able to maintain the appropriate environmental temperature.

Procedure

1. View sensor status.
   - From the CLI, run the **show sensors-fans** command.
   - If the output shows that all fans are running at 0 RPM, the fan module is not seated correctly in the appliance.
   - If the output shows that one or more fans are running at less than 1200 RPM, contact IBM Support.
2. View the fan module LED.
   - Amber single flash shows when power is first applied to the fan module.
   - Amber steady light indicates that the fan is operating at less than 1200 revolutions per minute (RPM) or there is a fault in the module.
   - No illumination when there is no power present or there is no problem.

What to do next

If the module is not seated correctly, remove and reinsert the module.

If you believe that the module must be replaced, contact IBM Support.

Troubleshooting power supply modules

How to troubleshoot the power supply module.

Procedure

1. View sensor status.
   - From the CLI, run the **show other-sensors** command.
2. View the power supply model LED.
   - Green steady light indicates that the module is connected to a power source.
   - Red steady light indicates that the module is not functioning within design specifications.
   - If not illuminated, there is no power to the module.
3. Remove the power cord from the power supply module. The appliance can operate with a single power supply module.

What to do next

If the module is not seated correctly, generally it is not locked in place. To ensure that the module is seated, remove and reinsert the module.

If the module has no AC power, ensure that the power cords are connected to the power supply and to a working AC power outlet.

If you believe that the module must be replaced, contact IBM Support.

Troubleshooting disk drive modules

How to troubleshoot the disk drive module.

Procedure

1. View RAID status.
   - From the CLI, run the **show raid-physical-drive** command.

   If the state shows `Unconfigured Bad`, the hard disk drive is damaged and must be replaced.
2. Contact IBM Support to replace the disk drive module.

Troubleshooting the appliance

You can use the **test hardware** command and the diagnostic self-test to troubleshoot your appliance.

When you can connect to the CLI, use the **test hardware** command to troubleshoot your appliance.

When you cannot connect to the CLI, use the boot-time diagnostic self-test to troubleshoot your appliance.

Removing or replacing the appliance or parts

The appliance parts can be removed or replaced under certain conditions.

The appliance includes two of three types of replacement parts: Tier 2 customer replaceable unit (CRU) and field replaceable unit (FRU). Following is a list of the three types of replacement part:

Tier 1 CRU
: Replacement of a Tier 1 CRU is your responsibility. If an IBM representative installs a Tier 1 CRU at your request, you are charged for the installation.

Tier 2 CRU
: Replacement of a Tier 2 CRU can be completed by you or an IBM representative for no charge if still under warranty. If installed by an IBM representative after your warranty expires, you are charged for the installation.

FRU
: Replacement of a FRU must be performed by an IBM representative only.

For information about the terms of warranty, see the *IBM Statement of Limited Warranty* document in the *Resource Kit*.

Removal and replacement guidelines

Read this information before you remove or replace a component.

- Review the guidelines for handling static-sensitive devices and the safety statements. This information helps you work safely.
- Observe good housekeeping in the area where you are working. Place removed parts in a safe place.
- You do not have to disconnect the appliance from the power supply to install or replace a hot-swap module if directed to do so.
- Ensure that enough properly grounded electrical outlets exist for the appliance.
- Have a medium Phillips screwdriver available.
- Component colors:
  - Orange
    - Orange on a component indicates that the component can be hot-swapped. You can remove or install the component while the appliance is running.

Orange can also indicate touch points on hot-swap components. See the instructions for removing or installing a specific hot-swap component for other procedures that you might have to complete before you remove or install the component.

– Blue
  - Blue on a component indicates touch points. You can grip touch points to remove or install the appliance, open, or close a latch, or for other purposes.

Guidelines for handling static-sensitive devices:

Read these guidelines before you handle static-sensitive devices.

Attention: Static electricity can damage the chassis and other electronic devices. To avoid damage, keep static-sensitive devices in their static-protective packages until you are ready to install them.

To reduce the possibility of electrostatic discharge, observe the following precautions:

- Limit your movement. Movement can cause static electricity to build up around you.
- The use of a grounding system improves safety. Wear an electrostatic-discharge wrist strap, if one is available.
- Handle the device carefully, holding it by its edges or its frame.
- Do not touch solder joints, pins, or bare circuitry.
- Do not leave the device where others can handle and damage it.
- While the device is still in its static-protective package, touch it to an unpainted metal part of the chassis or rack for at least 2 seconds. Touching the chassis drains static electricity from the package and from your body.
- Remove the device from its package and install it immediately without setting down the device. If it is necessary to set down the device, put it back into its static-protective package.
- Take extra care when you handle devices during times of cold weather. Indoor heating reduces ambient humidity and increases the conditions that cause static electricity to accumulate.

Returning an appliance or part:

If you are instructed to return an appliance or component, follow all packaging instructions and use the packaging materials that are provided for shipping.

Note: You might be charged for the replacement appliance or part if IBM does not receive the defective appliance or part within a reasonable amount of time. Contact IBM support with any questions.

Parts listing

The IBM MQ Appliance includes Tier 2 CRU parts and FRU parts.

For information about the terms of warranty, see the *IBM Statement of Limited Warranty* document on the *Resource Kit*.

CRU parts list - M2001:

The Ethernet modules, solid state disk drive modules, fan modules, power supply modules, and power cords are Tier 2 CRU parts.

Replacement of a Tier 2 CRU can be completed by you or an IBM representative for no charge if still under warranty. If installed by an IBM representative after your warranty expires, you are charged for the installation.

The following figure shows the CRU parts on the front and rear of the appliance.

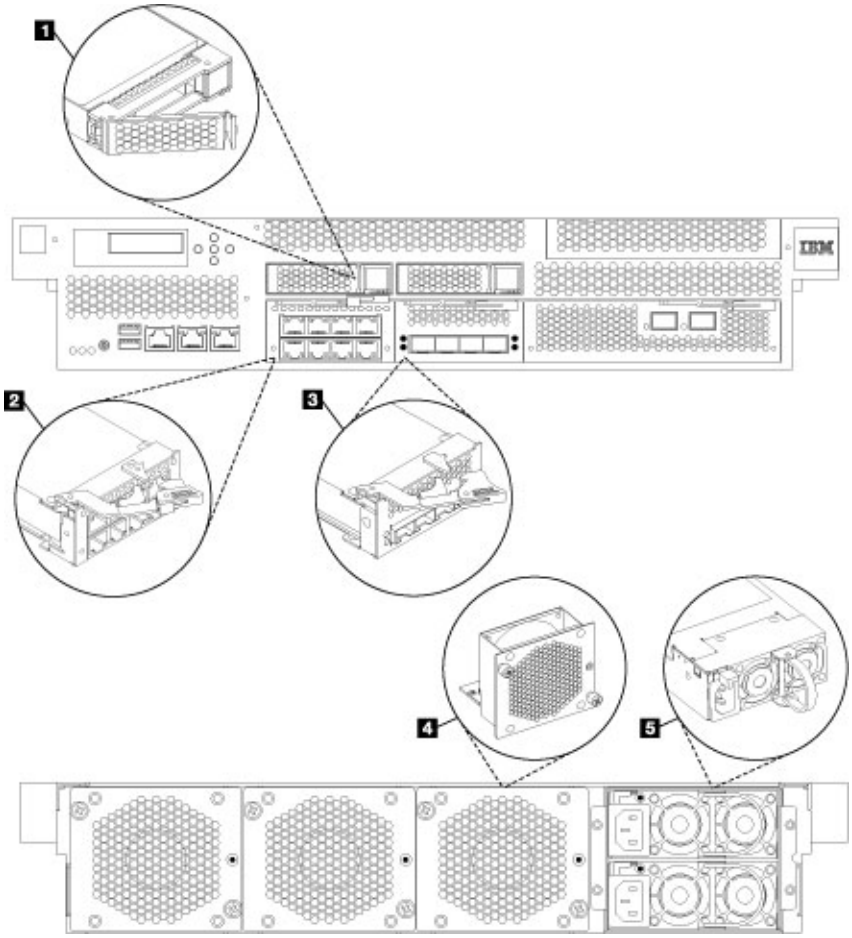

Figure 24. CRU part locations for the 8436 appliance.

The labels in this figure correspond to the following CRU components:

Table 9. Part numbers for the IBM MQ Appliance.

Label	Description	Tier 2 CRU part number
1	Solid state disk drive module assembly (complete)	00VM462
2	1 Gb Ethernet module with 8 ports for RJ45 interface	00VM052
3	10 Gb Ethernet module with 4 ports for SFP+ interface	00VM463
4	Fan module	97Y1290
5	Power supply module	97Y0440
The following CRU parts are not shown in the figure.		

Table 9. Part numbers for the IBM MQ Appliance. (continued)

Label	Description	Tier 2 CRU part number
-	DE-9 to RJ45 serial console cable	46N5656
-	USB to RJ45 serial console cable	97Y0517
-	Rail kit to mount the appliance into the rack.	60Y0328
-	Cat5e Ethernet cable x 2	01AF038
-	SFP+ direct attach copper wire Ethernet cable	90Y9432
-	SFP+ SR transceiver	46N5592

CRU parts list - M2000:

The Ethernet modules, hard disk drive modules, fan modules, power supply modules, and power cords are Tier 2 CRU parts.

Replacement of a Tier 2 CRU can be completed by you or an IBM representative for no charge if still under warranty. If installed by an IBM representative after your warranty expires, you are charged for the installation.

The following figure shows the CRU parts on the front and rear of the appliance.

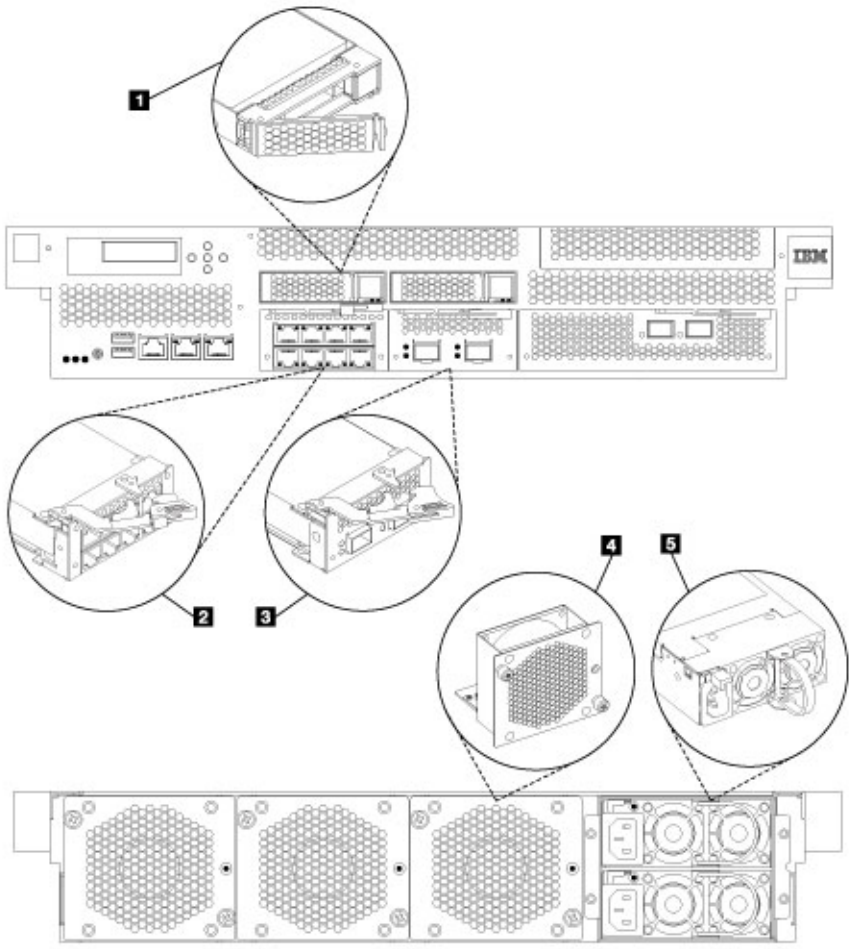

Figure 25. CRU part locations for the 8436 appliance.

The labels in this figure correspond to the following CRU components:

Table 10. Part numbers for the IBM MQ Appliance.

Label	Description	Tier 2 CRU part number
1	Hard disk drive module assembly (complete)	00VM039
2	1 Gb Ethernet module with 8 ports for RJ45 interface	00VM052
3	10 Gb Ethernet module with 2 ports for SFP+ interface	00VM312
4	Fan module	97Y1290
5	Power supply module	97Y0440
The following CRU parts are not shown in the figure.		
-	DE-9 to RJ45 serial console cable	46N5656
-	USB to RJ45 serial console cable	97Y0517
-	Rail kit to mount the appliance into the rack.	60Y0328
	SFP+ direct attach copper wire Ethernet cable	90Y9432
-	Cat5e Ethernet cable x 2	01AF038
-	SFP+ SR transceiver	46N5592

FRU parts listing:

FRU parts must be replaced only by an IBM representative.

The following table lists the FRU parts that are in the appliance.

Table 11. FRU part numbers for the appliance

Description	Part number
Shipping box	00VM076
2U chassis - 8436-54X	00VM631
2U chassis - 8436-55X	00VM675
16 GB DDR3 DIMM	00VM040
16 GB eUSB flash drive	00VM049
CMOS Button Cell battery	00RY543
CPU - Intel IvyBridge E5-2680-V2	00Y2786
RAID controller card and cache module	00VM235
RAID power backup capacitor and cable	00VM236
Emulex Fibre Channel card with carrier assembly	00VM053

Power cords

When you receive your appliance, the shipping carton contains power cords for rack mounted appliances.

To maintain warranty or service contracts, you must use IBM parts for power cords and rack cable cords.

Replacement of a Tier 2 CRU can be completed by you or an IBM representative for no charge if still under warranty. If installed by an IBM representative after your warranty expires, you are charged for the installation.

Table 12. Power cords and cords

Country	Tier 2 CRU part number	Description
Argentina	39M5068	2.8m, 10A/220V, C13 to IRAM 2073
Australia / New Zealand	39M5102	2.8m, 10A/250V, C13 to AS/NZ 3112
Brazil	39M5240	2.8m, 10A/125V, C13 to NBR 14136
Chile	39M5165	2.8m, 220 - 240V, C13 to CEI 23-16
China	39M5206	2.8m, 10A/250V, C13 to GB2099.1
Denmark	39M5130	2.8m, 10A/250V, C13 to DK2-5a
Europe	39M5123	2.8m, 10A/250V, C13 to CEE 7/7
India	39M5226	2.8m, 10A/250V, C13 IS 6538
Israel	39M5172	2.8m, 10A/250V, C13 to SI 32
Italy	39M5165	2.8m, 220 - 240V, C13 to CEI 23-16
Japan	39M5186	2.8m, 12A/240V, C13 to JIS C-8303
Japan	39M5199	2.8m, 12A/100V, C13 to JIS C-8303
Korea	39M5219	2.8m, 12A/250V, C13 to KSC 8305
South Africa	39M5144	2.8m, 10A/250V, C13 to SANS 164

Table 12. Power cords and cords (continued)

Country	Tier 2 CRU part number	Description
Switzerland	39M5158	2.8m, 10A/250V, C13 to SEV 1011-S24507
Taiwan	39M5247	2.8m, 10A/125V, C13 to CNS 10917-3
Taiwan	39M5254	2.8m, 10A/250V, C13 to CNS 10917-3
United Kingdom	39M5151	2.8m, 10A/250V, C13 to BS 1363/A
United States	39M5081	2.8m, 10A/125V, C13 to NEMA 5-15P
United States	39M5095	2.8m, 10A/250V, C13 to NEMA 6-15P
Rack power cords (all countries)	39M5377	2.8m, 10A/125-250 VAC, IEC 320 C13 to IEC 320 C14

Turning off the appliance

When the appliance must be turned off, use this procedure to turn off power to the appliance.

About this task

DANGER

Hazardous voltage, current, or energy levels are present inside. Do not open any cover or barrier. (L001)

Procedure

1. Save the changes from the running configuration to the startup configuration.

   From the IBM MQ Appliance web UI
   Click **Save Configuration**.

   From the CLI
   Use the **`write memory`** command.
2. Run the **`shutdown halt`** command to shut down the appliance.
3. Complete a graceful shutdown by pressing the power button at the front of the chassis.

What to do next

Verify that the power LED at the front of the appliance is not illuminated. To remove all power from the system, the power cords must be unplugged from both power supply units.

Removing and replacing CRU parts

Use this hardware maintenance procedure to remove and replace a CRU part when directed by IBM Support.

About this task

Replacement of a Tier 2 CRU can be completed by you or an IBM representative for no charge if still under warranty. If installed by an IBM representative after your warranty expires, you are charged for the installation.

Procedure

- “Replacing a fan module”
- “Replacing a power supply module” on page 77
- “Replacing a solid state disk drive module - M2001 appliances” on page 78
- “Replacing an Ethernet module” on page 85

Replacing a fan module:

How to replace a failed fan module.

Before you begin

You must have part 97Y1290 available.

You must turn off the appliance and replace a fan module when directed by IBM Support.

About this task

When one or more fan modules are not working, turn off the appliance as soon as possible to avoid overheating. The remaining fans might not be able to maintain the appropriate environmental temperature.

DANGER

Hazardous voltage, current, or energy levels are present inside. Do not open any cover or barrier. (L001)

DANGER

Rack-mounted devices are not to be used as shelves or work spaces. (L002)

DANGER

Multiple power cords. The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. (L003)

Procedure

1. If the appliance is not turned off, complete a graceful shutdown by pressing the power button at the front of the appliance. Wait until the power LED is no longer illuminated to indicate that the appliance power is turned off.
2. Unplug all network cables and power cords.
3. Remove the fan module.

   The following figure shows the numbered components that are mentioned in the steps.

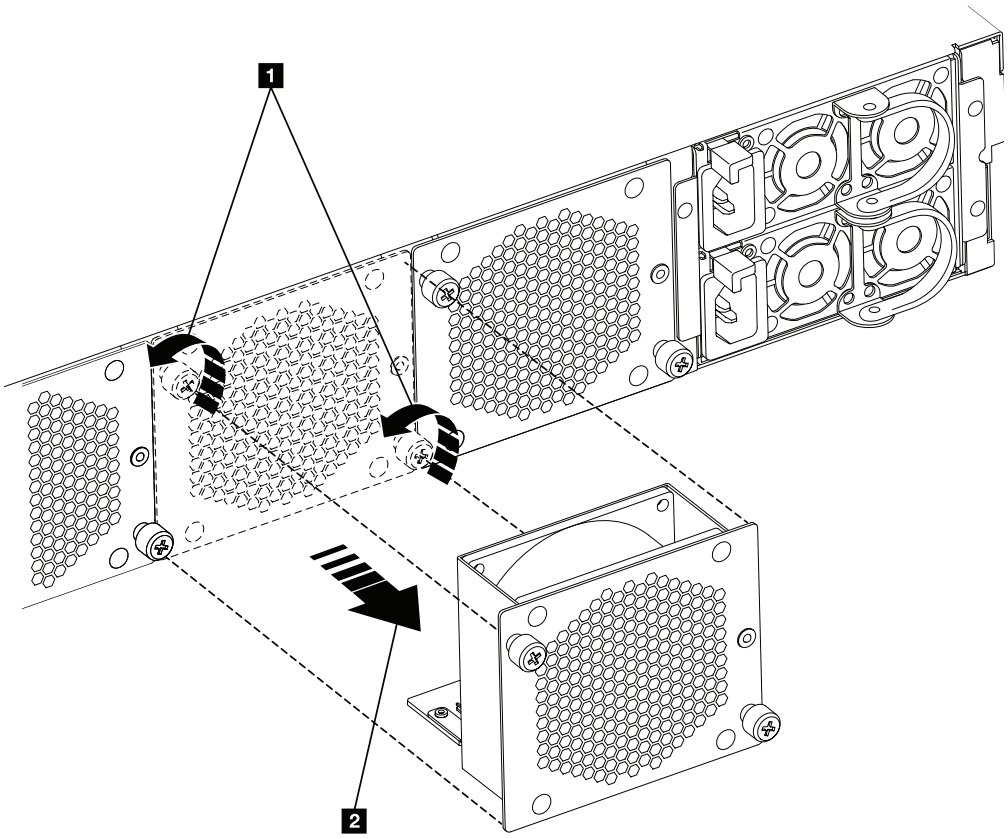

Figure 26. Removing a fan module

a. Unscrew the two thumbscrews on the fan module until they twist without resistance **1**. The fan module thumbscrews are designed to remain attached to the fan module.
b. Pull the fan module to remove it from the appliance **2**.

4. Set the faulty module aside.

   Attention: Ensure that the gold connectors at the rear of the module do not come into contact with your hands or with the packing material as you unpack the replacement module. Avoid damaging the gold connectors against the appliance as you insert the replacement module.
5. Unpack the replacement module.
6. Carefully align the replacement module, and insert until the module face is flush with the rear panel.
7. Tighten the thumbscrews on the fan module.
8. Plug in all power cords.
9. Turn on the appliance by pressing the power button.
10. After you replace the fan module, confirm that the new module is working by verifying that the following are true.
    a. The fan module LED is not illuminated.
    b. The fault LED at the front of the appliance is not illuminated.

What to do next

After you verify that the replacement module is working, return the failed part to IBM.

Replacing a power supply module:

Use this procedure to replace a power supply module.

Before you begin

You must have purchased a power supply module. The part number of a power supply module is 97Y0440.

About this task

There are two hot-swap power supplies in the rear of the appliance. You need to replace a power supply module as soon as possible when directed by IBM Support or if any of the following situations occur.

- When the appliance generates a critical or warning message to indicate which power supply module is in a failure state.
- When the LED on one of the power supply modules is illuminated red.
- The amber fault LED at the front of the appliance is illuminated when a hardware fault is detected.

DANGER

Hazardous voltage, current, or energy levels are present inside. Do not open any cover or barrier. (L001)

DANGER

Rack-mounted devices are not to be used as shelves or work spaces. (L002)

Procedure

1. Unplug the power cord of the failed module.
2. Remove the power supply module.

   The following figure shows the numbered components that are mentioned in the steps.

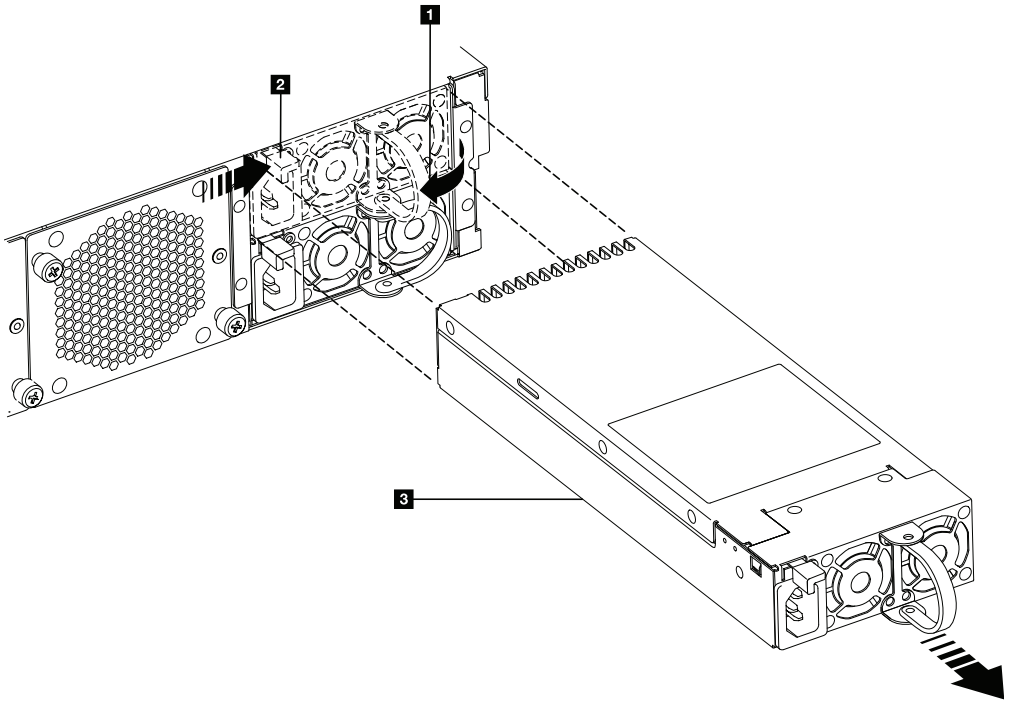

Figure 27. Removing a power supply module.

a. Rotate, then firmly grip the handle **1** of the failed module.
b. Push the orange release latch **2** toward the handle **1** and hold in this position.
c. Pull the failed module from the appliance **3** .

3. When fully removed from the appliance, set aside the failed module.

   Attention: Ensure that the gold connectors at the rear of the module do not come into contact with your hands or with the packing material as you unpack the replacement module. Avoid damage to the gold connectors as you insert the replacement module.
4. Unpack the replacement module.
5. Replace the module.
   a. Carefully align the replacement module with the open space in the appliance.
   b. Completely insert the module until the release latch clicks into place.
   c. Pull the handle to ensure that the module is secure.
6. Plug in the power cord to the replaced module.
7. Verify that the new module is working.
   a. The power supply LED is illuminated green.
   b. The fault LED is not illuminated.

What to do next

After you verify that the replacement module is working, return the failed part to IBM.

Replacing a solid state disk drive module - M2001 appliances:

How to replace a solid state disk drive module.

Before you begin

The part number of a solid state disk drive module is 00VM461.

The solid state disk drive modules are not hot-swappable. Hot swapping the modules causes your system to crash, and might damage your appliance. You must turn off the appliance before you replace the solid state disk drive module.

About this task

You need to replace a solid state disk drive module when the solid state disk state is `Unconfigured Bad` or if directed by IBM Support.

DANGER

When you work on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- **Connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.**
- **Do not open or service any power supply assembly.**
- **Do not connect or disconnect any cables or install, maintain, or reconfigure this product during an electrical storm.**
- **The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.**
- **Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.**
- **Connect any equipment that is attached to this product to properly wired outlets.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following procedures when you install, move, or open covers on this product or attached devices.**

To disconnect:

1. **Turn off everything (unless instructed otherwise).**
2. **Remove the power cords from the outlets.**
3. **Remove the signal cables from the connectors.**
4. **Remove all cables from devices.**

To connect:

1. **Turn off everything (unless instructed otherwise).**
2. **Attach all cables to devices.**
3. **Attach the signal cables to the connectors.**
4. **Attach the power cords to the outlets.**
5. **Turn on the devices.**

(D005)

Procedure

1. If the appliance is not turned off, complete a graceful shutdown by pressing the power button at the front of the appliance. The green power LED turning off indicates that the appliance is powered off.

   The following figure shows the numbered components that are mentioned in the steps.

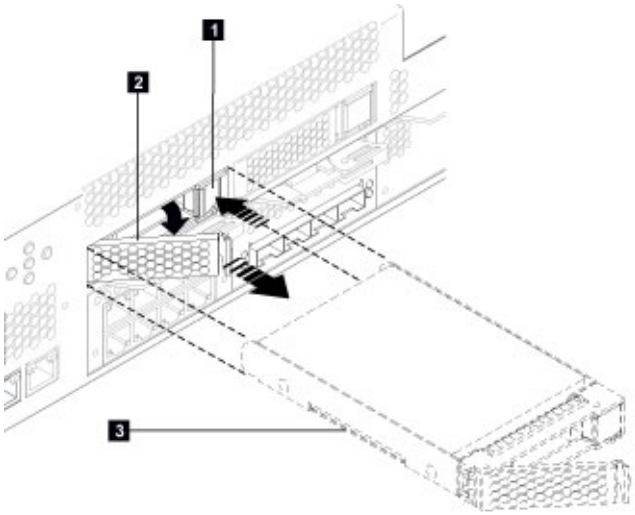

Figure 28. Removing a solid state disk drive module.

2. Press the locking arm release latch **1** and the locking arm is released.
3. To unlock the module, rotate the locking arm approximately 40 degrees by pulling out **2**.
4. To remove the module, pull the module out of the appliance **3**.
5. Set aside the failed module.

   Attention: Ensure that the gold connectors at the rear of the module do not come into contact with your hands or with the packing material as you unpack the replacement module. Avoid damaging the gold connectors against the appliance as you insert the replacement module.
6. Unpack the replacement module.
7. Carefully align the module, and insert into the opening until the module is seated.
8. Push the locking arm towards the appliance until the release latch clicks into place.
9. Connect all network cables and power cords.
10. Turn on the appliance by pressing the power button that is on the front of the appliance.
11. Verify that the power LED is illuminated steady green.
12. Verify that the new module is working.
    a. The solid state disk drive activity LED illuminates steady green.
    b. The solid state disk state is not `Unconfigured Bad`.

What to do next

After you verify that the replacement module is working, return the failed part to IBM.

Replacing a hard disk drive module - M2000 appliances:

How to replace the hard disk drive module.

Before you begin

The part number of a hard disk drive module is 00VM039.

The hard disk drive modules are not hot-swappable. Hot swapping the modules causes your system to crash, and might damage your appliance. You must turn off the appliance before you replace the hard disk drive module.

About this task

You need to replace a hard disk drive module when the hard disk state is `Unconfigured Bad` or if directed by IBM Support.

DANGER

When you work on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- **Connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.**
- **Do not open or service any power supply assembly.**
- **Do not connect or disconnect any cables or install, maintain, or reconfigure this product during an electrical storm.**
- **The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.**
- **Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.**
- **Connect any equipment that is attached to this product to properly wired outlets.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following procedures when you install, move, or open covers on this product or attached devices.**

To disconnect:

1. **Turn off everything (unless instructed otherwise).**
2. **Remove the power cords from the outlets.**
3. **Remove the signal cables from the connectors.**
4. **Remove all cables from devices.**

To connect:

1. **Turn off everything (unless instructed otherwise).**
2. **Attach all cables to devices.**
3. **Attach the signal cables to the connectors.**
4. **Attach the power cords to the outlets.**
5. **Turn on the devices.**

(D005)

Procedure

1. If the appliance is not turned off, complete a graceful shutdown by pressing the power button at the front of the appliance. The green power LED turning off indicates that the appliance is powered off.

   The following figure shows the numbered components that are mentioned in the steps.

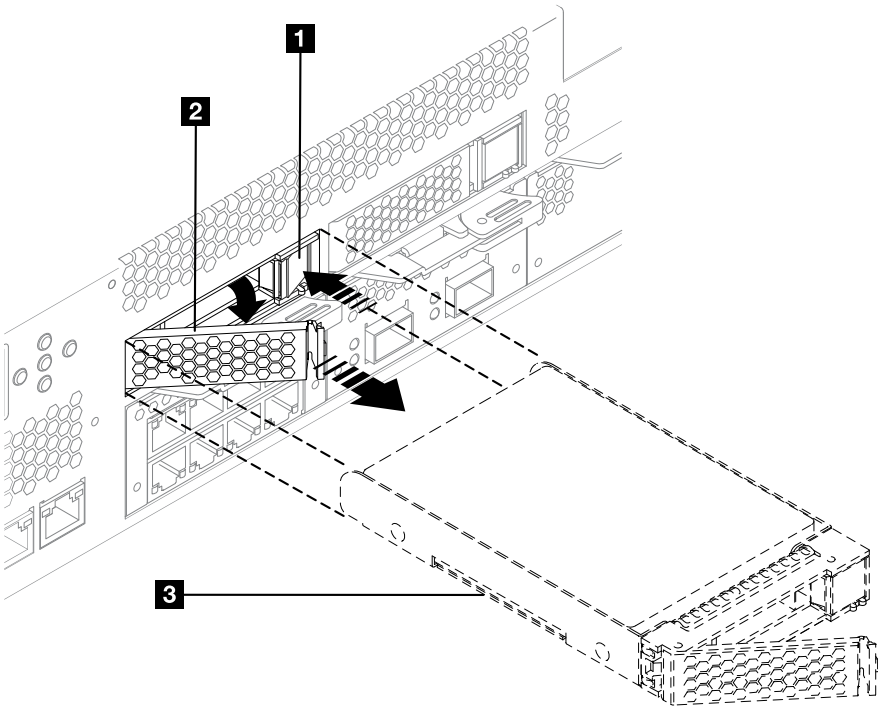

Figure 29. Removing a hard disk drive module.

2. Press the locking arm release latch **1** and the locking arm is released.
3. To unlock the module, rotate the locking arm approximately 40 degrees by pulling out **2** .
4. To remove the module, pull the module out of the appliance **3** .
5. Set aside the failed module.

   Attention: Ensure that the gold connectors at the rear of the module do not come into contact with your hands or with the packing material as you unpack the replacement module. Avoid damaging the gold connectors against the appliance as you insert the replacement module.
6. Unpack the replacement module.
7. Carefully align the module, and insert into the opening until the module is seated.
8. Push the locking arm towards the appliance until the release latch clicks into place.
9. Connect all network cables and power cords.
10. Turn on the appliance by pressing the power button that is on the front of the appliance.
11. Verify that the power LED is illuminated steady green.
12. Verify that the new module is working.
    a. The hard disk drive activity LED illuminates steady green.
    b. The hard disk state is not `Unconfigured Bad`.

What to do next

After you verify that the replacement module is working, return the failed part to IBM.

Replacing an Ethernet module:

The procedure to replace an Ethernet module.

Before you begin

You must have purchased an Ethernet module.

- The part number of the 1 GB Ethernet module is 00VM052.
- The part number of the 4 x 10 GB Ethernet module is 00VM455.

You must turn off the appliance before you replace the Ethernet module. When you disconnect network cables from the appliance, be sure to label each so that you can connect them in the proper location.

About this task

Removal instructions are the same for both modules.

You can replace an Ethernet module if you have a problem with your module or if directed by IBM Support if the following situation occurs.

- You are unable to connect to the network even though the cable is plugged in.
- If the output from the **test hardware** command includes `Expected number of interfaces: x - found y`.
- When you use listing, all the Ethernet ports in the module are not included in the list:
  - From the IBM MQ Appliance web UI select **Manage Appliance** > **Network** > **Ethernet Interface** > **..**
  - From the CLI, use the **show interface** command.

DANGER

When you work on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:
- **Connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.**
- **Do not open or service any power supply assembly.**
- **Do not connect or disconnect any cables or install, maintain, or reconfigure this product during an electrical storm.**
- **The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.**
- **Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.**
- **Connect any equipment that is attached to this product to properly wired outlets.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following procedures when you install, move, or open covers on this product or attached devices.**

To disconnect:
1. **Turn off everything (unless instructed otherwise).**
2. **Remove the power cords from the outlets.**
3. **Remove the signal cables from the connectors.**
4. **Remove all cables from devices.**

To connect:
1. **Turn off everything (unless instructed otherwise).**
2. **Attach all cables to devices.**
3. **Attach the signal cables to the connectors.**
4. **Attach the power cords to the outlets.**
5. **Turn on the devices.**

(D005)

DANGER

Multiple power cords. The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. (L003)

Procedure

1. If the appliance is not turned off, complete a graceful shutdown by pressing the power button at the front of the appliance. When the power LED is no longer illuminated, the appliance is powered off.

   The following figure shows the numbered components that are mentioned in the steps.

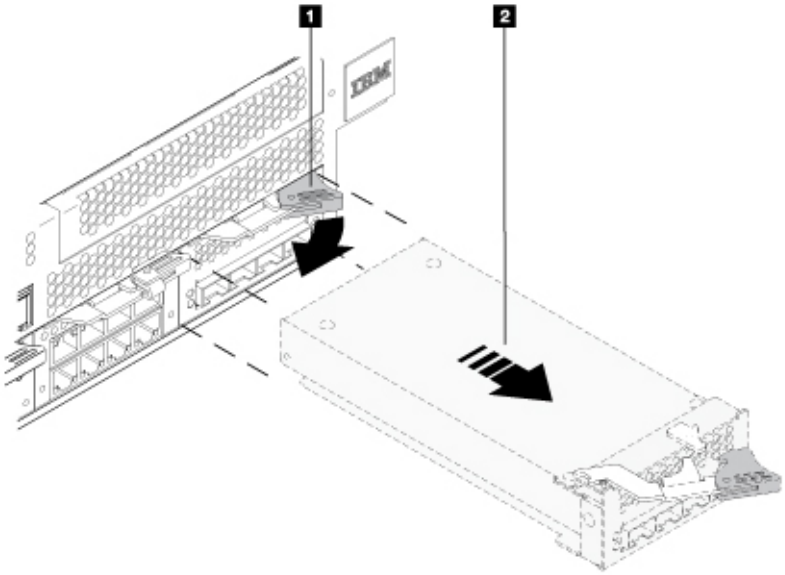

Figure 30. Removing the 10 Gb Ethernet module (M2001 model).

2. Grasp the blue latch **1** rotate slightly and pull outward.
3. Pull the module out of the appliance **2** with care to support the module weight as it exits.
4. Set aside the Ethernet module.

   Attention: Ensure that the gold connectors at the rear of the module do not come into contact with your hands or with the packing material as you unpack the replacement module. Avoid damaging the gold connectors against the chassis as you insert the replacement module.
5. Unpack the replacement module.
6. Carefully align the module, and insert into the appliance.
7. Push the Ethernet module forward until the module is securely in place.
8. Push the blue latch back in place to lock the module.
9. Turn on the appliance by pressing the power button at the front of the appliance and verify that the power LED is illuminated steady green.
10. After you replace the module, verify that the new module is working.
    a. You can connect to the network after you plug in the cable and the activity LED is illuminated.
    b. The fault LED light is not illuminated.

What to do next

After you verify that the replacement module is working, return the failed part to IBM.

Removing the appliance from the rack

After you install the appliance in the rack, you generally remove it only to move it to another position in the rack.

About this task

DANGER

Rack-mounted devices are not to be used as shelves or work spaces. (L002)

CAUTION:

The weight of this part or unit is 18 - 32 kg (39.7 - 70.5 lb). It takes two persons to safely lift this part or unit. (C009)

Procedure

1. If the appliance is not turned off, press the power button on the front of the chassis. The power LED is no longer illuminated when the power is turned off.
2. Unplug all power cords from the appliance.

   The following figure shows the numbered components that are mentioned in the steps.

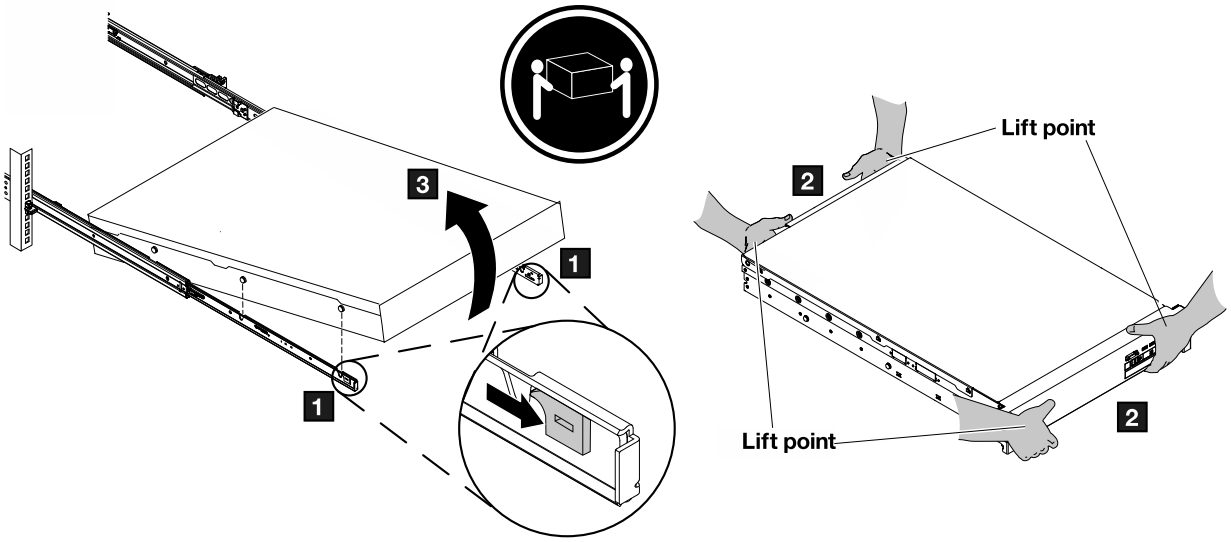

Figure 31. Unlatching and rotating the front of the appliance.

3. Separate the appliance from the rails.
   a. Pull the locking levers **1** forward.
   b. Make sure that two people support the front and the rear of the appliance at lifting points **2** .
   c. Lift the front of the appliance up slightly **3** to clear the nailhead from the slot.
   d. Unlatch and lift the front of the appliance.

e.

4. Lift the appliance directly from the rails.
   a. After the front nailheads clear the latches, lift the rear of the appliance to make the appliance level.
   b. Lift the appliance directly out of the rack from Lift points 1 , and 2 .
5. Place the appliance on a sturdy, clean surface.
6. Slide the rails back in the rack.

Removing the batteries

How to remove the battery and capacitor for end-of-life recycling.

About this task

DANGER

Multiple power cords. The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. (L003)

DANGER

Improper disposal or incineration of batteries or capacitors can cause life-threatening injury.

The Type 8436 appliance does not have any internal user serviceable parts. Any battery or capacitor is to be accessed and removed only by trained personnel. These instructions apply only to end-of-life recycling procedures.

Procedure

1. Turn off the appliance and disconnect all power cords and external cables from the appliance.
2. Remove the cover of the appliance.
3. Locate the CMOS battery on the system board next. The battery is next to the RAM slots.

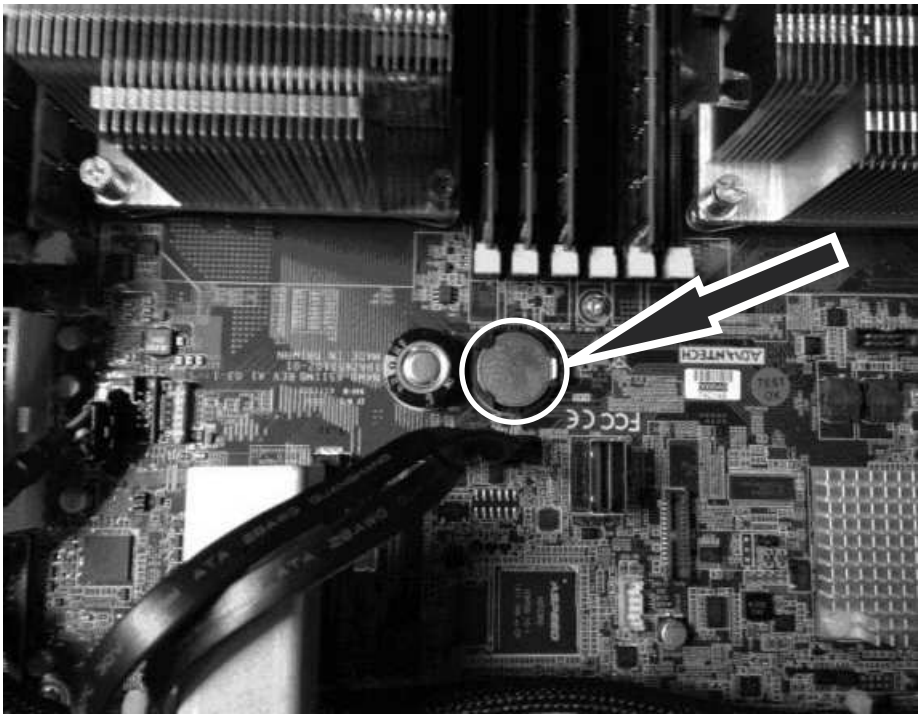

4. Remove the battery with your fingers to release and lift the battery from the connector.

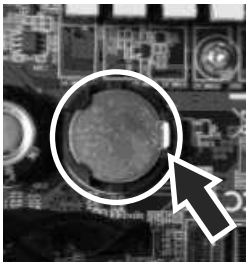

5. Locate the RAID capacitor inside the chassis.

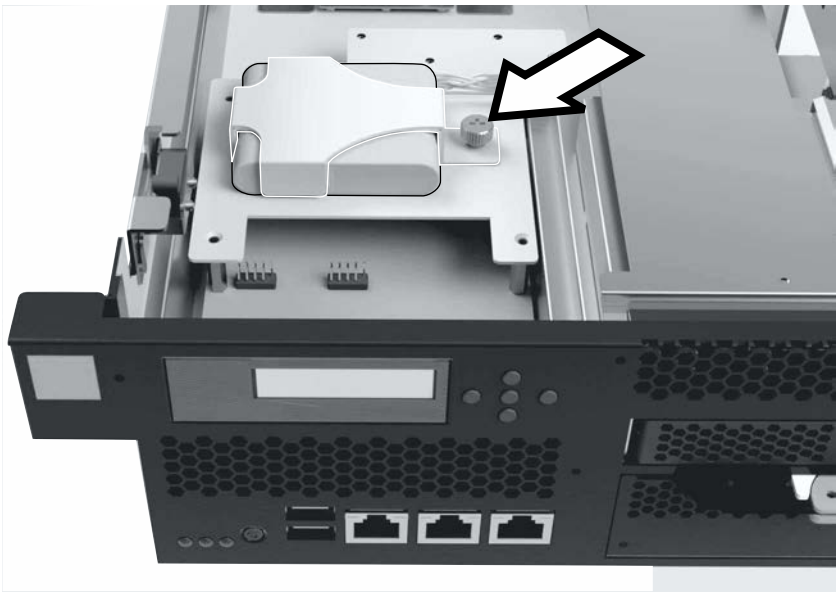

6. Loosen the indicated capacitor cover retention screw to remove the capacitor cover.
7. Disconnect the RAID capacitor power connector and remove the capacitor from the appliance.

What to do next

Dispose of batteries and capacitors as required by your local ordinances or regulations.

Chapter 4. Upgrading and downgrading

To upgrade your IBM MQ Appliance, you must install the latest level of firmware on the appliance.

New function, security updates, and maintenance fixes for the IBM MQ Appliance are made available through firmware releases. Additional maintenance through iFixes are made available, as necessary, on the most recent firmware level release.

You can also downgrade your appliance, if required, either by reverting to the previous level of firmware or by installing a specific firmware version.

Fixes are cumulative, so you should always download the most recent firmware that is available on the IBM Fix Central website.

All queue manager information persists through upgrades and rollbacks.

Both appliances in a high availability pair should have the same firmware level. Appliances can operate at different levels to allow time to upgrade the appliances separately, but you should avoid configuring HA queue managers during this period.

Both appliances in a disaster recovery configuration should have the same firmware level. Appliances can operate at different levels to allow time to upgrade the appliances separately, but you should avoid configuring DR queue managers during this period.

The standard rules of applying upgrades to IBM MQ apply when firmware is upgraded and queue managers restarted. If the new firmware moves the IBM MQ installation to a new command level, any queue managers that are started at the new command level are no longer able to start under a lower command level, even following a firmware rollback.

If you have a high availability (HA) configuration, you must ensure that you update both appliances to the new command level at the same time. Queue managers running at one command level are not able to run on an appliance with a different command level.

If you have a disaster recovery (DR) configuration, you must update the recovery appliance to the new command at the same time that you update the main appliance. This ensures that a queue manager that has run on the main appliance can be started on the recovery appliance if required.

Installing a new level of firmware

You upgrade the IBM MQ Appliance by downloading a new version of the firmware and installing it.

About this task

You download the new firmware image from the IBM Fix Central website, and then copy the image to the appliance. You then reboot the appliance, and use the new image.

© Copyright IBM Corp. 2015, 2016

You can use this process to upgrade to a new level of firmware, or to downgrade to a specific, earlier version of the firmware.

If the appliance that you are upgrading is part of a high availability configuration, then you pause the first appliance, then upgrade and resume the appliance. You then pause, upgrade, then resume the second appliance. See "Suspending an appliance from an HA group for maintenance" on page 209 for guidance.

Procedure

1. Use a computer with web access to download the required image from IBM Fix Central. This website is a repository for all available and supported firmware images for IBM MQ Appliances. The fixes are cumulative, so always choose the most recent image.
2. Back up your IBM MQ Appliance as described in "Back up and restore" on page 198.
3. Restart the appliance as described in "Restarting the appliance" on page 198.
4. Ensure that all the queue managers are stopped.
5. Copy the firmware image from the computer that you downloaded it to to the `image:` location on the appliance.

   To copy the image by using the command line:

   a. Connect to the command line of the appliance as described in "Command line access" on page 97.
   b. Log in to the appliance as an administrator.
   c. Type `config` to enter configuration mode.
   d. Type `flash` to enter the correct mode for firmware upgrade.
   e. Copy the file by typing the following command:

      ```
      copy scp://username@ipaddress/[/]directorypath/firmware_file image:
      ```

   To copy the image by using the IBM MQ Appliance web UI:

   a. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
   b. Click **Actions** for the image folder.
   c. Select **Upload files** from the **Actions** menu.
   d. Click **Browse**, and browse for the location of the image file on your local system.
   e. Click **Upload** to upload the image file to the image directory on the appliance.
6. Restart the appliance with the new image by typing the following command:

   ```
   boot image accept-license firmware_file
   ```

   Where *firmware_file* is the name of the file that contains the new firmware image. Type the file name without the `image:` prefix.
7. When the appliance restarts, verify that the firmware image is upgraded by entering the following command:

   ```
   show version
   ```

Reverting to the previous level of firmware

You can, if instructed, revert to a previous level of firmware.

About this task

When you upgrade the IBM MQ Appliance firmware by using the `boot image` command, the appliance retains current configuration data. This feature is used to restore the appliance to a known, stable state if required.

- The previous firmware image and associated configuration data is the secondary installation.
- The newly installed firmware image and associated configuration data is the primary installation.

When you switch between firmware images, the switch can take some time. During this switch operation, do not power off or restart the appliance.

When you run the **`boot switch`** command, all system configuration reverts to its state at initial upgrade (the **`boot image`** step). Any changes made since the firmware update are lost. Losses include, for example, changes to network configuration or user definitions. If any changes made immediately following the upgrade are the cause of problems (and trigger the rollback), this state should be rectified by the **`boot switch`** operation. However, data associated with queue managers is not modified when the boot switch takes place, so, for example, messages are not lost in the rollback, and security certificates and authority records retain any modifications.

You cannot perform this operation on an appliance that is part of a high availability group.

Procedure

1. Connect to the command line of the appliance as described in "Command line access" on page 97.
2. Log in to the appliance as an administrator.
3. Ensure that all queue managers are stopped.
4. Type `config` to enter configuration mode.
5. Type `flash` to enter the correct mode for firmware roll-back.
6. Restart the appliance with the original image by typing the following command:

   ```
   boot switch
   ```

Downgrading

You can downgrade your appliance, if required, either by reverting to the previous level of firmware or by installing a specific firmware version.

To revert to the previous level of firmware, follow the procedure described in "Reverting to the previous level of firmware" on page 94.

To install a specific firmware version, follow the procedure described in "Installing a new level of firmware" on page 93.

If you encounter a startup configuration error when you restart the appliance after downgrading the firmware, complete the following steps:

1. Save the current configuration by completing the procedure described in "Backing up or saving the appliance configuration" on page 199.

2. Restart the appliance by following the procedure described in "Restarting the appliance" on page 198.

Suspending an appliance from an HA group for maintenance

When you want to suspend an appliance from a high availability group, for example, to carry out maintenance on the appliance, you perform a managed failover. This procedure transfers all the workload to the remaining appliance in the group.

To achieve the managed failover, you put the appliance that you want to temporarily remove from the group into standby mode. You then resume the appliance after the maintenance is complete.

Note: While you have one appliance in standby mode, your queue managers can run only on the remaining appliance. You should take care to avoid any outage on the second appliance.

You use this technique when you update the firmware on the appliances in your high availability group, for example to apply a fix pack. In this situation, you suspend the first appliance, update the firmware, and then resume it. You can then suspend the other appliance, upgrade the firmware, and then resume it.

Chapter 5. Configuring

You can configure the IBM MQ Appliance settings by using the IBM MQ Appliance web UI or by using the command line.

Command line access

You must access the command line on the IBM MQ Appliance before you can enter appliance administrative or IBM MQ commands.

You can access the command line in the following ways:

- Directly connect to the serial port. For details, see "Connecting to the serial port."
- Enable SSH and establish a remote connection through the SSH service. For details, see "Configuring the SSH service."

To use the IBM MQ control commands you must enter the IBM MQ administration mode by entering the command `mqcli` on the command line. You can exit the IBM MQ administration mode by entering the command `exit`.

Connecting to the serial port

The serial port is hard-wired to a command line administration shell.

Before you connect to the serial port, ensure that the configuration for the terminal or PC is for standard 9600 8N1 and no flow control operation. 8N1 is a notation for a serial configuration in asynchronous mode, where there are eight (8) data bits, no (N) parity bit, and one (1) stop bit.

To make the serial connection, use the appropriate cable to connect the terminal or PC to the console connector on the appliance. You require a serial-to RJ45 converter cable or USB-to-RJ45 converter cable to connect.

When properly connected, the terminal or PC prompts for credentials for a locally defined user. With a serial connection, the following restrictions apply:

- Authentication does not use RBM. Therefore, only locally defined users can log in to the serial port. If a non-local user attempts to log in, the log contains the following message:

  ```
  auth error 0x81000034  User 'user001' failed to log in.
  ```

- There can be only one active serial connection at a time.
- In addition to being local, the user must be the admin account or a privileged user.

Configuring the SSH service

By default, the SSH service is disabled. When enabled, the SSH service binds to the defined local IP-address-port combination.

Without an explicit local address, the SSH service attempts to bind to the management Ethernet interface. If the management Ethernet interface is not defined, the SSH service binds to all configured interfaces.

© Copyright IBM Corp. 2015, 2016

Be sure to define an explicit IP address to isolate management traffic from application data traffic.

If any of the Ethernet interfaces on the appliance are connected to the internet, or a similar open access network, you might want to prevent access to the SSH service from those interfaces. By restricting the Ethernet interface that can be used to access the SSH service, you can ensure that the service can be accessed only from an internal network. This restriction makes your environment more secure.

To establish an SSH session

Although many servers use password authentication for SSH login, the IBM MQ Appliance requires an interactive process to protect credentials during the SSL handshake. The IBM MQ Appliance initiates a secure channel and provides for an encrypted login process.

As a side-effect of the initial connection, and depending on your SSH client, you might see the extraneous "login as:" prompt. To bypass, press Enter.

The screen shows a warning about unauthorized access and the prompt for the login credentials:

```
login as:
Unauthorized access prohibited.
login:
```

Configuring the SSH service by using the IBM MQ Appliance web UI

You can configure the SSH service by using the **SSH Service** pane in the IBM MQ Appliance web UI.

Procedure

1. In the IBM MQ Appliance web UI, select **Manage Appliance** > **Network Access** > **SSH Service**.
2. Select **Enable administrative state**.
3. In the **Local address** field, enter the local address the appliance monitors for incoming SSH requests.
4. Optional: In the **Port number** field, change the port on which the appliance monitors for incoming SSH requests.
5. Optional: Click **Edit** next to the **Access control list** field to modify the SSH ACL.
6. Click **Apply** to save the changes to the running configuration.
7. Optional: Click **Save Configuration** to save the changes to the persisted configuration.

Configuring the SSH service by using the command line

You can configure the SSH service by using the **config** and **ssh** commands on the command line.

Procedure

1. Ensure that you are not in the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

   The prompt displays mqa#.

2. Enter the configuration mode by entering the following command:
   `config`
3. Specify the IP address and port that the SSH service listens on by entering the following command:
   `ssh IPAddress:Port`
   where:

   IPAddress
   Specifies the IP address of the Ethernet interface that you want to use to access the SSH service.

   port Specifies the port number that you want to use to access the SSH service.
4. Exit the configuration mode by entering the following command:
   `exit`

Configuring the IBM MQ Appliance web UI

The IBM MQ Appliance web UI can be used to administer the IBM MQ Appliance.

The IBM MQ Appliance web UI can be accessed by using any of the supported browsers. The browser must have JavaScript enabled. Currently supported browsers are the latest versions of Firefox or Chrome, or Internet Explorer 11.

Note: If you are using Internet Explorer, you must ensure that you are running in standards mode, not compatibility mode. To change the mode, press F12. Check the **Document Mode:** menu, and the **Browser Mode:** menu, and make any necessary changes.

You can connect to the IBM MQ Appliance web UI by entering the following URL:

Note: This URL uses the default port value. If you changed the port value, replace the `9090` section of the URL with your port number.

`https://IP_Address:9090`

Where:

IP_Address
Specifies the IP address of the management Ethernet interface.

You can determine the IP address of the management Ethernet interface by using the **`show int`** command.

Changing the IBM MQ Appliance web UI IP address and port

System administrators can change the IP address and port values that are used to connect to the IBM MQ Appliance web UI by using the **`config`** and **`web-mgmt`** commands on the command line.

About this task

If any of the Ethernet interfaces on the appliance are connected to the internet, or a similar open access network, you might want to prevent access to the IBM MQ Appliance web UI from those interfaces. By restricting the Ethernet interface that can be used to access the IBM MQ Appliance web UI, you can ensure that the IBM MQ Appliance web UI can be accessed only from an internal network. This

restriction makes your environment more secure.

Procedure

1. Ensure that you are not in the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

   The prompt displays mqa#.
2. Enter the configuration mode by entering the following command:

   ```
   config
   ```
3. Enter the web management configuration mode by entering the following command:

   ```
   web-mgmt
   ```
4. Enable the web management service by entering the following command:

   ```
   admin-state enabled
   ```
5. Configure which IP address is used to access the IBM MQ Appliance web UI:
   - Allow access to the IBM MQ Appliance web UI through all Ethernet interfaces and all IP addresses by entering the following command:

     ```
     local-address 0.0.0.0:port
     ```

     Where:

     port Specifies the port number that you want to use to access the IBM MQ Appliance web UI.
   - Allow access to the IBM MQ Appliance web UI through only a specific Ethernet interface and IP address by entering the following command:

     ```
     local-address IPAddress:port
     ```

     Where:

     IPAddress
     Specifies the IP address of the Ethernet interface that you want to use to access the IBM MQ Appliance web UI.

     port Specifies the port number that you want to use to access the IBM MQ Appliance web UI.
6. Exit the web management configuration mode by entering the following command:

   ```
   exit
   ```
7. Exit the configuration mode by entering the following command:

   ```
   exit
   ```

Configuring certificates for IBM MQ Appliance web UI

You can configure the IBM MQ Appliance web UI to use certificates that you supply.

About this task

You use the appliance command line interface to configure the IBM MQ Appliance web UI to use your certificates.

To set up secure communication between a browser and the IBM MQ Appliance web UI and to handle certificates, you create an SSL server profile on the appliance. You import the required certificates and key file to the appliance, and create definition objects for them. The definition objects are used when you create

an ID credentials (idcred) object for the appliance. The idcred is in turn used when you configure the SSL server profile. Finally, the SSL server profile is associated with your web management profile.

If you want to configure client validation, you import the certificates of the clients that are going to be allowed to connect. You then create definition objects for the certificates, which are used when you create a validation credential (valcred) object. The valcred object is in turn used when you configure the SSL server profile.

The example in this topic assumes that you have a signed certificate for the appliance. When making certificate requests for an appliance, the CN part of the distinguished name must be the URL that you type to reach the web UI. For example, `myappliance1.ourcompany.com`. If you want to set up the profile to validate connecting clients, you also require the relevant client certificates.

By default the web management service listens on all of the appliance ports (`local address` set to 0.0.0.0). You can, however, configure the service so that it listens on an IP address or host alias of a specific port (and so limit access to the web UI - see "Changing the IBM MQ Appliance web UI IP address and port" on page 99).

Procedure

- To upload certificates to your appliance:
  1. Ensure that you have the following items:
     - A private key to access the appliance certificate.
     - The appliance certificate.
     - Client certificates (optional).
  2. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
  3. Log in as a user in the administrators group.
  4. Type the following command to enter configuration mode:
     ```
     config
     ```
  5. Upload the key and certificates to the appliance by using the `copy` command, for example:
     ```
     copy scp://username@otherserver//home/username/myappliance1key.pem cert:
     copy scp://username@otherserver//home/username/myappliance1.cer cert:
     copy scp://username@otherserver//home/username/client1.cer cert:
     copy scp://username@otherserver//home/username/client2.cer cert:
     copy scp://username@otherserver//home/username/client3.cer cert:
     ```

  You can also copy the certificates to your appliance by using the IBM MQ Appliance web UI, see "Uploading certificates to the appliance" on page 277.
- To create definition objects for the appliance certificate and key:
  1. From configuration mode, type `crypto` to enter crypto configuration mode.
  2. Create a crypto key definition for the private key that is used for generating the appliance certificate:
     ```
     key key_alias cert:///keyfile
     ```

     For example:
     ```
     key WebUiKey01 cert:///myappliance1key.pem
     ```
  3. Create a crypto certificate definition for the appliance:
     ```
     certificate cert_alias cert:///certfile
     ```

     For example:

```
certificate WebUiCert01 cert:///myappliance1.cer
```

4. Create a crypto credential definition for the appliance:

   ```
   idcred credential_name key_alias cert_alias
   ```

   For example:

   ```
   idcred WebUiCred01 WebUiKey01 WebUiCert01
   ```

- To create a crypto valcred definition for validating clients (this is optional):
  1. From the crypto configuration mode, create a certificate definition object for each of the client certificates that you have imported:

     ```
     certificate cert_alias cert:///certfile
     ```

     For example:

     ```
     certificate WebUiClientCert01 cert:///client1.cer
     certificate WebUiClientCert02 cert:///client2.cer
     certificate WebUiClientCert03 cert:///client3.cer
     ```

  2. Create a crypto valcred definition, specifying the certificate definitions for the client certificates:

     ```
     valcred valcred_name
     certificate cert_alias
     ```

     Repeat the **certificate** command to specify the certificate definition for every client certificate that you have uploaded. For example:

     ```
     valcred WebUIvalcred01
     certificate WebUIClientCert01
     certificate WebUIClientCert02
     certificate WebUIClientCert03
     ```

- To create an SSL server profile for the appliance:
  1. From the crypto configuration mode, enter the following commands:

     ```
     ssl-server SSL_Svr_Profile_name
     admin-state enabled
     idcred IDCred_name
     protocols TLSv1d2
     ```

     If you are specifying client validation, also enter:

     ```
     valcred ValCred_name
     request-client-auth on
     require-client-auth on
     send-client-auth-ca-list on
     ```

     For example:

     ```
     ssl-server myappliance1
     admin-state enabled
     idcred WebUiCred01
     protocols TLSv1d2
     valcred WebUIvalcred01
     request-client-auth on
     require-client-auth on
     send-client-auth-ca-list on
     ```

- To save all the changes you have made in crypto configuration mode:
  1. Type `exit` to leave crypto configuration mode.
  2. Type `write mem` to save your configuration changes.
- To associate the SSL server profile with the web UI:
  1. From configuration mode, type `web-mgmt` to enter web management configuration mode.
  2. Enter the following command:

```
ssl-server SSL_Svr_Profile_name
```

For example:

```
ssl-server myappliance1
```

- To save your web management configuration:
  1. Type `exit` to leave web-mgmt configuration mode.
  2. Type `write mem` to save your configuration changes.
  3. Type `exit` again to leave configuration mode.

Configuring the appliance

You can configure the appliance either by using the IBM MQ Appliance web UI or by using the command line.

Ethernet interfaces

You must configure the Ethernet interfaces on your IBM MQ Appliance to enable communications.

The IBM MQ Appliance provides two management ports and up to twelve Ethernet ports:

- The management ports are `mgt0` and `mgt1`. You can use either port for management traffic.
- The Ethernet ports are in two modules.
  - The left module has eight 1-gigabit ports. Port-numbering starts with 10, which is `eth10`.
  - For M2000 models, the right module has two 10-gigabit ports. Port-numbering starts with 20, which is `eth20`.
  - For M2001 models, the right module has four 10-gigabit ports. Port-numbering starts with 20, which is `eth20`.

Note:

1. By default, the appliance blocks non-management traffic when at least one network interface has an invalid configuration. When this situation occurs, the appliance supports only management traffic over SSH and web management interfaces. Until you correct the problem, the appliance cannot accept and process client requests.
2. You must not disable Ethernet links that are used for high availability or disaster recovery configurations. If you do disable such links, high availability operation or disaster recovery operation are no longer available, and you might have to set up the configuration again.
3. You must not change IP addresses for Ethernet links that are used for high availability or disaster recovery configurations. If you do change IP addresses, high availability operation or disaster recovery operation are no longer available, and you might have to set up the configuration again. See “Changing IP addresses in high availability configurations” on page 129 and “Changing IP addresses in disaster recovery configurations” on page 146.
4. You must not aggregate Ethernet links that are used for high availability or disaster recovery configurations. If you do aggregate such links, high availability operation or disaster recovery operation are no longer available, and you might have to set up the configuration again.

Configuring Ethernet interfaces by using the IBM MQ Appliance web UI

You can configure the Ethernet interfaces by using the IBM MQ Appliance web UI.

About this task

You use the Manage Appliance part of the IBM MQ Appliance web UI to configure the Ethernet interfaces on your IBM MQ Appliance. The IBM MQ Appliance web UI itself contains detailed help on the fields that you need to configure. Here we describe the major steps that you need to complete.

Note: There are restrictions on changing the IP addresses of the Ethernet ports used by high availability and disaster recovery configurations. See "Changing IP addresses in high availability configurations" on page 129 and "Changing IP addresses in disaster recovery configurations" on page 146.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage** > **Network** and then click **Ethernet interface**.
3. Click **New** to define a new configuration.
4. Specify a name for the configuration. You must specify the name of the Ethernet interface that you want to configure. The following names are available:
   - eth10
   - eth11
   - eth12
   - eth13
   - eth14
   - eth15
   - eth16
   - eth17
   - eth20
   - eth21
   - eth22 (M2001 appliances only)
   - eth23 (M2001 appliances only)
   - mgt0
   - mgt1
5. Define the basic configuration for that interface. Here you specify whether the interface is active or not, and whether the interface uses static addressing, or an autconfiguration scheme. If you choose either of the autoconfiguration schemes, the appliance ignores the remaining configuration data about the physical interface.
6. Define IP addressing. Here you define the primary IP address and netmask for the configuration. You can optionally specify one or more secondary addresses.
7. Define IP routing. Here you define default gateways for IPv4 and IPv6 IP addresses. You can optionally set up a routing table that defines static routes.
8. Optionally define advanced options that modify the way that the Ethernet interface works within the network.
9. Click **Apply** to save the changes to the running configuration.

10. Click **Save Configuration** to save the changes to the named configuration.

Configuring Ethernet interfaces by using the command line

You can configure the Ethernet interfaces by using the command interface.

About this task

To configure an Ethernet interface from the command line, you enter Ethernet configuration mode, specifying the interface to configure, and enter the required Ethernet commands.

Note: There are restrictions on changing the IP addresses of the Ethernet ports used by high availability and disaster recovery configurations. See "Changing IP addresses in high availability configurations" on page 129 and "Changing IP addresses in disaster recovery configurations" on page 146.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure an Ethernet interface:

   ```
   ethernet name
   ```

   Where name identifies the interface that you want to configure and has one of the following values:
   - eth10
   - eth11
   - eth12
   - eth13
   - eth14
   - eth15
   - eth16
   - eth17
   - eth20
   - eth21
   - eth22 (M2001 appliances only)
   - eth23 (M2001 appliances only)
   - mgt0
   - mgt1
4. Use the following Ethernet interface commands to configure the interface:

Command	Description
"**admin-state** on page 437	Sets the administrative state for the configuration.
"**flow-control** on page 495	Sets the flow control mode of the Ethernet interface.
"**ip-address** on page 496	Assigns the primary network address for the Ethernet interface.
"**ip-config-mode** on page 497	Identifies the configuration mode for the Ethernet interface.
"**ip-route** on page 498	Manages static routes in the routing table for the Ethernet interface.
"**ip-secondary-address** on page 499	Manages secondary network addresses for the Ethernet interface.

Command	Description
"**ipv4-default-gateway**" on page 499	Designates the default IPv4 gateway for the Ethernet interface.
"**ipv6-dadtransmits**" on page 500	Sets the number of IPv6 duplication address detection attempts for the Ethernet interface.
"**ipv6-default-gateway**" on page 500	Designates the default IPv6 gateway for the Ethernet interface.
"**ipv6-nd-retransmit-timer**" on page 501	Sets the interval between IPv6 neighbor discovery attempts for the Ethernet interface.
"**link**" on page 5011	Indicates whether the Ethernet interface is part of an aggregate interface.
"**mac-address**" on page 502	Changes the MAC address for the Ethernet interface.
"**mode**" on page 503	Sets the interface speed and direction.
"**mtu**" on page 503	Sets the maximum transmission unit of the Ethernet interface.

5. Use the following commands to control the interface, if required:

Command	Description
"**disable-ethernet-hardware-offload**" on page 496	Manages the temporary disabling of hardware offload.
"**packet-capture**" on page 504	Manages a packet-capture for the Ethernet interface session.

6. After you configure the Ethernet interface, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Example

The following commands configure interface `eth10` to use DHCP:

```
mqa# config
Global configuration mode
mqa(config)# ethernet eth10
Modify Ethernet Interface configuration

mqa(config ethernet eth10)# ip-config-mode dhcp
mqa(config ethernet eth10)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

VLAN interfaces

You can configure the VLAN interfaces on your IBM MQ Appliance.

A VLAN interface allows multiple logical LANs to coexist on the same Ethernet segment. VLAN packets are identified by the IEEE 802.1Q protocol. You can define multiple VLAN interfaces on a single parent interface. The parent interface can be an Ethernet interface or a link aggregation interface.

VLAN interfaces are not supported for links used for high availability configurations or disaster recovery configurations.

Note: By default, the appliance blocks non-management traffic when at least one network interface has an invalid configuration. When this situation occurs, the appliance supports only management traffic over SSH and web management interfaces. Until you correct the problem, the appliance cannot accept and process client requests.

Configuring VLAN interfaces by using the IBM MQ Appliance web UI

You can configure the VLAN interfaces by using the IBM MQ Appliance web UI.

About this task

You use the Manage Appliance part of the IBM MQ Appliance web UI to configure VLAN interfaces on your IBM MQ Appliance. The IBM MQ Appliance web UI itself contains detailed help on the fields that you need to configure. Here we describe the major steps that you need to complete.

Normally, IP addresses are fixed configuration. Use Dynamic Host Configuration Protocol (DHCP) or Stateless Address Autoconfiguration (SLAAC) only when you have a reason to obtain addresses dynamically. When you select an autoconfiguration method, the appliance ignores configuration data about the physical interface.

You must configure the parent interface before you configure the VLAN interface.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage** > **Network** and then click **VLAN interface**.
3. Click **New** to define a new configuration.
4. Specify a name for the configuration.
5. Define the basic configuration for that interface. Here you specify whether the interface is active or not, and whether the interface uses static addressing, or an autconfiguration scheme, and whether the parent interface is Ethernet or link aggregation. If you choose either of the autoconfiguration schemes, the appliance ignores the remaining configuration data about the physical interface.
6. Specify the parent interface of the VLAN. Specify the name of either the Ethernet or link aggregation configuration.
7. Define the identifier for the VLAN you are configuring, and the priority level for outbound VLAN headers for packets.
8. Define IP addressing. Here you optionally define the primary IP address and one or more secondary addresses.
9. Define IP routing. Here you optionally define default gateways for IPv4 and IPv6 IP addresses. You can set up a routing table that defines static routes.
10. Optionally enable standby control and define a standby configuration.
11. Optionally define advanced options that modify the way that the VLAN interface works within the network.
12. Click **Apply** to save the changes to the running configuration.
13. Click **Save Configuration** to save the changes to the named configuration.

Configuring VLAN interfaces by using the command line

You can configure VLAN interfaces by using the command interface.

About this task

To configure a VLAN interface from the command line, you enter VLAN configuration mode, specifying the interface to configure, and enter the required VLAN commands.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure a VLAN interface:

   `vlan` *name*

   Where name identifies the interface that you want to configure. The name can have a maximum of 128 characters. The following characters are valid:
   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name that consists of a single period, or including two periods together, is not permitted)
4. Use the following VLAN interface commands to configure the interface:

Command	Description
"**admin-state**" on page 437	Sets the administrative state for the configuration.
"**ethernet-interface**" on page 614	Sets the Ethernet interface to provide connectivity.
"**identifier**" on page 615	Sets the VLAN identifier.
"**ip-address**" on page 615	Assigns the primary network address for the VLAN interface.
"**ip-config-mode**" on page 616	Identifies the configuration mode for the VLAN interface.
"**ip-route**" on page 616	Manages static routes in the routing table for the VLAN interface.
"**ip-secondary-address**" on page 617	Manages secondary network addresses for the VLAN interface.
"**ipv4-default-gateway**" on page 618	Designates the default IPv4 gateway for the VLAN interface.
"**ipv6-dadtransmits**" on page 618	Sets the number of IPv6 duplication address detection attempts for the VLAN interface.
"**ipv6-default-gateway**" on page 619	Designates the default IPv6 gateway for the VLAN interface.
"**ipv6-nd-retransmit-timer**" on page 619	Sets the interval between IPv6 neighbor discovery attempts for the VLAN interface.
"**link-aggregation-interface**" on page 620l	Indicates whether the VLAN interface is part of an aggregate interface.
"**mtu**" on page 620	Sets the maximum transmission unit of the VLAN interface.
"**outbound-priority**" on page 621	Sets the priority value in outbound packets.
"**over**" on page 621	Sets the parent interface type.

5. Use the following command to control the interface, if required:

Command	Description
"**packet-capture**" on page 622	Manages a packet-capture for the VLAN interface session.

6. After you configure the VLAN interface, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Link aggregation interfaces

You can configure the link aggregation interfaces on your IBM MQ Appliance.

A link aggregation interface combines multiple Ethernet ports. When combined and used in parallel, the aggregate interface increases link speed beyond a single Ethernet port. Because the aggregate interface combines Ethernet ports, redundancy is increased to provide higher availability.

You must configure the Ethernet interfaces, and enable them for link aggregation, before you create a link aggregation interface.

Link aggregation is not supported for links used for high availability configurations or disaster recovery configurations.

Note: By default, the appliance blocks non-management traffic when at least one network interface has an invalid configuration. When this situation occurs, the appliance supports only management traffic over SSH and web management interfaces. Until you correct the problem, the appliance cannot accept and process client requests.

Configuring link aggregation interfaces by using the IBM MQ Appliance web UI

You can configure a link aggregate interface by using the IBM MQ Appliance web UI.

About this task

You use the Manage Appliance part of the IBM MQ Appliance web UI to configure a link aggregate interface on your IBM MQ Appliance. The IBM MQ Appliance web UI itself contains detailed help on the fields that you need to configure. Here we describe the major steps that you need to complete.

Note: You must configure the Ethernet interfaces that you want to aggregate, and enable them for link aggregation, before you create the link aggregation interface.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage** > **Network** and then click **Link Aggregation interface**.
3. Click **New** to define a new configuration.
4. Specify a name for the link aggregation interface.
5. Define the basic configuration for that interface. Here you specify whether the interface is active or not, and whether the interface uses static addressing, or an autoconfiguration scheme. If you choose either of the autoconfiguration schemes, the appliance ignores the remaining configuration data about the physical interface.
6. Specify the aggregation mode for the interface.
7. Specify the Ethernet interfaces to be included in the aggregation.
8. Define IP addressing. Here you define the primary IP address and netmask for the configuration. You can optionally specify one or more secondary addresses.
9. Define IP routing. Here you define default gateways for IPv4 and IPv6 IP addresses. You can optionally set up a routing table that defines static routes.

10. Optionally define advanced options that modify the way that the link aggregation interface works within the network.
11. Click **Apply** to save the changes to the running configuration.
12. Click **Save Configuration** to save the changes to the named configuration.

Configuring link aggregation interfaces by using the command line

You can configure the link aggregation interfaces by using the command line interface.

About this task

To configure a link aggregation interface from the command line, you enter link aggregation configuration mode, specifying the interface to configure, and enter the required link aggregation commands.

Note: You must configure the Ethernet interfaces that you want to aggregate, and enable them for link aggregation, before you create the link aggregation interface.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure a link aggregation interface:

   ```
   link-aggregation name
   ```

   Where *name* identifies the interface that you want to configure. The name can have a maximum of 128 characters. The following characters are valid:
   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name that consists of a single period, or including two periods together, is not permitted)
4. Use the following link aggregation interface commands to configure the interface:

Command	Description
"**admin-state**" on page 437	Sets the administrative state for the configuration.
"**ip-address**" on page 542	Assigns the primary network address for the link aggregation interface.
"**ip-config-mode**" on page 543	Identifies the configuration mode for the link aggregation interface.
"**ip-route**" on page 544	Manages static routes in the routing table for the link aggregation interface.
"**ip-secondary-address**" on page 545	Manages secondary network addresses for the link aggregation interface.
"**ipv4-default-gateway**" on page 545	Designates the default IPv4 gateway for the link aggregation interface.
"**ipv6-dadtransmits**" on page 546	Sets the number of IPv6 duplication address detection attempts for the link aggregation interface.

Command	Description
"**ipv6-default-gateway**" on page 546	Designates the default IPv6 gateway for the link aggregation interface.
"**ipv6-nd-retransmit-timer**" on page 547	Sets the interval between IPv6 neighbor discovery attempts for the link aggregation interface.
"**lacp-hash**" on page 547	Sets which hash function the LACP aggregation uses to determine the interface for outbound packets.
"**lacp-select**" on page 548	Sets the algorithm for the LACP selection policy.
"**link**" on page 548	Specifies which Ethernet interfaces are part of the aggregate interface.
"**type**" on page 550	Defines the mode to use for link aggregation.

5. Use the following command to control the interface, if required:

Command	Description
"**packet-capture**" on page 549	Manages a packet-capture for the aggregate interface session.

6. After you configure the link aggregation interface, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

IPMI Settings

You can define an IPMI channel, and IPMI users on the IBM MQ Appliance.

An Intelligent Platform Management Interface (IPMI) LAN channel enables access to the Baseboard Management Controller (BMC) on the appliance over a LAN.

An IPMI user can create, change, or delete user authentication records in the BMC. Authentication records allow users to communicate with IPMI protocols over external channels, such as an IPMI LAN channel. There can be eight IPMI users on the appliance.

Configuring IPMI LAN channels by using the command line interface

You can configure an IPMI LAN channel by using the command line interface.

About this task

You configure the appliance Ethernet port `mgt0` to provide an IPMI LAN channel, which can include serial over LAN (SOL) access.

To configure an IPMI user from the command line, you enter IPMI channel mode, and enter the required IPMI channel commands.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type `ipmi-lan-channel` to enter IPMI LAN channel mode.
4. Use the following commands to configure the channel:

Command	Description
"**admin-state**" on page 437	Sets the administrative state for the configuration.

Command	Description
"**allowed-user**" on page 538	Specifies the users allowed to use the channel.
"**ip address**" on page 539	Sets IP addresses with subnet mask for the IPMI LAN channel.
"**ip default-gateway**" on page 540	Sets the default gateway for the IPMI LAN channel.
"**maximum-channel-privilege-level**" on page 540	Sets the maximum privilege level for users.
"**sol-enabled**" on page 541	Indicates whether to support serial over LAN.
"**sol-required-user-privilege-level**" on page 541	Sets the privilege level for serial over LAN.

5. After you configure the channel, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Configuring IPMI users by using the command line interface

Define an IPMI user to create an identification record in the Baseboard Management Controller (BMC).

About this task

To configure an IPMI user from the command line, you enter IPMI user mode, and enter the required IPMI user commands. Users are identified by an integer. Possible identifiers are 3 to 10 inclusive.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type `ipmi-user` to enter IPMI user mode.
4. Enter the following command to define an IPMI user:
   `user-id` *`identifier`*

   Where *identifier* is an integer in the range 3 to 10 inclusive.
5. Enter the following command to define a password for the user:
   `password` *`password`*

   Where *password* specifies a password that must be between 8 and 16 characters in length.
6. After you configure the IPMI user, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

DNS settings

DNS setting define the DNS servers to contact to resolve host names to IP addresses.

The primary behavior to configure for host name resolution consists of the following definitions:

- Which domains to search for a match when a host name without a domain qualifier is submitted. The appliance attempts to resolve a host name with any domain in the domain name table. The host name is resolved to the first found match.
- Which DNS servers to contact and their contact order.

- The load-balancing algorithm to contact name servers:
  - **First alive**. This algorithm maintains a list of servers and forwards a new connection to the next server on the list.
  - **Round robin**. This algorithm uses the concept of a primary server and one or more backup servers. When the primary server is available, all connections are forwarded to this server. When the primary server is unavailable, connections are forwarded to backup servers. The primary server is the first server in the list.

The results from DNS resolution requests are cached to improve performance. When a name server responds with an IP address, the response includes its time to live (TTL) in the cache. The appliance uses the value from the DNS response or 10 seconds, whichever is greater. If the name server responds that a host name has no associated IP address, the appliance caches the negative response for 30 seconds.

Configuring DNS settings by using the IBM MQ Appliance web UI

You use the IBM MQ Appliance web UI to configure the DNS settings for the IBM MQ Appliance.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage** > **Network** and then click **DNS service**.
3. Enable the DNS service.
4. Set whether you prefer IPv4 or IPv6 addresses.
5. Define the domains to search to match partial host name. Use the directional arrows to set the search order.
6. Define the name servers to contact for resolution. Use the directional arrows to set the contact order. For the first alive algorithm, you can set a maximum of three name servers.
7. Define host-address maps for static hosts. (The use of static hosts does not improve performance).
8. Select the load balancing algorithm.
   - For the first alive algorithm, set the maximum number of query attempts before an error is returned, and the time to wait for a response before querying the next name server in the list.
   - For the round robin algorithm, set the maximum number of query attempts on a per-server basis.
9. Click **Apply** to save the changes to the running configuration.
10. Click **Save Configuration** to save the changes to the persisted configuration.

Configuring DNS settings by using the command line

You can configure the DNS settings by using the command interface.

About this task

To configure DNS settings from the command line, you enter DNS configuration mode and enter the required DNS commands.

Procedure

1. Connect to the IBM MQ Appliance as described in “Command line access” on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.

3. Type the following command to configure DNS settings:
   `dns`
4. Use the following DNS commands to configure the settings:

Command	Description
"**admin-state**" on page 437	Sets the administrative state for the configuration.
"**ip-preference**" on page 490	Sets the preferred IP version when the DNS provider publishes both versions of addresses.
"**load-balance**" on page 491	Sets the load-balancing algorithm that the appliance uses to resolve host names.
"**name-server**" on page 491	Manages local DNS providers.
"**retries**" on page 492	Sets the number of times that the appliance attempts a failed query.
"**search-domain**" on page 493	Manages domain-suffixes in the search table for non-qualified domain names.
"**static-host**" on page 493	Manages host-address maps.
"**timeout**" on page 494	Sets the time to wait before the next query attempt.

5. After you configure the DNS settings, enter `exit`
6. Enter `write memory` to save the updated configuration.

Clearing hosts from the DNS cache

The IBM MQ Appliance maintains a cache of DNS hosts. If the cache contains unwanted hosts, you must clear the cache.

About this task

You can clear hosts from the cache from the DNS Settings configuration.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage** > **Network** and then click **DNS service**.
3. Click **Actions** > **Flush DNS Cache**.

Configuring the locale, date, and time

Configure a connection to a network time protocol server to manage the locale, date, and time on the IBM MQ Appliance.

You can configure the date and time automatically by using network time protocol (NTP) servers to synchronize the server time with another server. Using NTP servers to synchronize the time ensures that the IBM MQ Appliance time is automatically updated, and is therefore always correct. It is important for the IBM MQ Appliance time to be correct as it is used in key information such as log files and monitoring data.

You can configure a connection to an NTP server by using the web UI or by using the command line interface.

Configuring NTP service settings by using the IBM MQ Appliance web UI

You can use the IBM MQ Appliance web UI to configure the NTP service settings for the IBM MQ Appliance.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage** > **Network** and then click **NTP service**.
3. Enter the host name or IP address of the NTP server. Click **Add** to specify multiple NTP servers. The servers are contacted in the order that you specify them.
4. In the **Refresh interval** field, specify the interval between clock synchronizations.
5. Click **Apply** to save the changes to the running configuration.
6. Click **Save Configuration** to save the changes to the persisted configuration.

Configuring NTP service settings by using the command line

You can configure the NTP service settings by using the command interface.

About this task

To configure NTP service settings from the command line, you enter NTP configuration mode and enter the required NTP commands.

After an NTP server is identified by the `remote-server` command, the appliance functions as a Simple Network Time Protocol (SNTP) client as described in RFC 2030. As an NTP client, the appliance issues time-of-day requests to the specified NTP server every 15 minutes.

The appliance supports one NTP server at a time. To designate a new NTP server, use the `no ntp-service` command to delete the current server, and then use the `remote-server` command to designate the new server.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure NTP service settings:
   ```
   ntp-service
   ```
4. Use the following NTP commands to configure the settings:

Command	Description
"**admin-state** on page 437	Sets the administrative state for the configuration.
"**refresh-interval** on page 489	Sets the interval between clock synchronizations.
"**remote-server** on page 490	Identifies an NTP server.

5. After you configure the NTP service settings, enter `exit`
6. Enter `write memory` to save the updated configuration.

Configuring the timezone by using the IBM MQ Appliance web UI

You use the IBM MQ Appliance web UI to configure the timezone settings for the IBM MQ Appliance.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage Appliance** > **Appliance** and then click **Time Settings**.
3. Select the timezone from the **Local time zone** list.

4. Click **Apply** to save the changes to the running configuration.
5. Click **Save Configuration** to save the changes to the persisted configuration.

Configuring timezone settings by using the command line

You can configure the timezone settings by using the command interface.

About this task

To configure timezone settings from the command line, you enter timezone mode and enter the required timezone commands.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure timezone settings:
   ```
   timezone
   ```
4. Use the following timezone commands to configure the settings:

Command	Description
"**admin-state**" on page 437	Sets the administrative state for the configuration.
"**custom**" on page 478	Specifies the name of a custom time zone.
"**daylight-name**" on page 479	Specifies the name of the time zone when in Daylight Saving Time. This name is appended to the time.
"**daylight-offset**" on page 479	Sets the offset, in hours, for Daylight Saving Time.
"**daylight-start-day**" on page 479	Specifies the day of the week when Daylight Saving Time starts.
"**daylight-start-hours**" on page 480	Specifies the hour in the day when Daylight Saving Time starts.
"**daylight-start-minutes**" on page 481	Specifies the minute in the hour when Daylight Saving Time starts.
"**daylight-start-month**" on page 481	Specifies the month of the year when Daylight Saving Time starts.
"**daylight-start-week**" on page 482	Specifies the instance of the day in the month when Daylight Saving Time starts.
"**daylight-stop-day**" on page 483	Specifies the day of the week when Daylight Saving Time stops.
"**daylight-stop-hours**" on page 483	Specifies the hour in the day when Daylight Saving Time ends.
"**daylight-stop-minutes**" on page 484	Specifies the minute in the hour when Daylight Saving Time ends.
"**daylight-stop-month**" on page 485	Specifies the month of the year when Daylight Saving Time ends.
"**daylight-stop-week**" on page 486	Specifies the instance of the day in the month when Daylight Saving Time ends.
"**direction**" on page 486	Specifies the direction, relative to Coordinated Universal Time, of the time zone.
"**name**" on page 487	Sets the name of the time zone. This name is appended to the time.
"**offset-hours**" on page 488	Specifies the offset in hours, relative to Coordinated Universal Time, of the time zone.
"**offset-minutes**" on page 489	Specifies the offset in minutes, relative to Coordinated Universal Time, of the time zone.

5. After you configure the timezone settings, enter `exit`. The setting of the value of the timezone becomes effective in the appliance when you issue the `exit` command. If you do not want to make the setting effective, type `cancel` instead.
6. Restart the appliance to ensure that the new timezone is used by all IBM MQ processes.

Configuring the locale by using the IBM MQ Appliance web UI

You use the IBM MQ Appliance web UI to configure the locale for the IBM MQ Appliance.

About this task

First, you enable the language that the locale uses, and then you select the locale.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage Appliance** > **Appliance** and then click **Language**.
3. Double-click the language that you want to enable to open its properties.
4. Select **Enable administrative state**, then click **Apply**.
5. Select **Manage Appliance** > **Appliance** and then click **System Settings**.
6. Select the locale from the **System Locale** list.
7. Click **Apply** to save the changes to the running configuration.
8. Click **Save Configuration** to save the changes to the persisted configuration.

Configuring the locale by using the command line

You can configure the locale by using the command interface.

About this task

First, you enable the language that the locale uses, and then you specify the locale.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to enable the language for the locale:

   ```
   language language_code
   admin-state enabled
   ```

   Where language code specifies the language that you want to enable and is one of the following codes:
   - de (German)
   - en (English)
   - es (Spanish)
   - fr (French)
   - it (Italian)
   - ja (Japanese)
   - ko (Korean)
   - pt_BR (Brazilian Portuguese)
   - ru (Russian)

- zh_CN (Simplified Chinese)
- zh_TW (Chinese, Taiwan)

4. Type the following command to enter system settings mode:
   `system`
5. Type the following command to specify the locale:
   `locale` *`language_code`*
6. After you configure the locale settings, enter `exit`
7. Enter `write memory` to save the updated configuration.

Configuring the appliance name

Configuring the appliance name helps you to identify the appliance on which you are operating.

You must configure the appliance name if the appliance is part of a high availability configuration or a disaster recovery configuration.

Configuring the appliance name by using the IBM MQ Appliance web UI

You can configure the appliance name by using the IBM MQ Appliance web UI.

About this task

You use the Manage Appliance part of the IBM MQ Appliance web UI to configure the appliance name on your IBM MQ Appliance. The IBM MQ Appliance web UI itself contains detailed help on the fields that you need to configure. Here we describe the major steps that you need to complete.

Note: There are restrictions on changing the names of appliances used by high availability and disaster recovery configurations. See "Changing appliance names in high availability configurations" on page 130 and "Changing appliance names in disaster recovery configurations" on page 146.

Procedure

1. Start the IBM MQ Appliance web UI, and click **Manage Appliance**.
2. Select **Manage Appliance** > **Appliance** and then click **System Settings**.
3. Specify the appliance name in the **Appliance name** field. Use letters and numbers in the name, avoid special characters. The name should not consist of numbers only.
4. Click **Apply** to save the changes to the running configuration.
5. Click **Save Configuration** to save the changes to the named configuration.

Configuring the appliance name by using the command line

You can configure the appliance name by using the command interface.

About this task

To configure the appliance name from the command line, you enter system settings mode and enter the required system name configuration command.

Note: There are restrictions on changing the names of appliances used by high availability and disaster recovery configurations. See "Changing appliance names in high availability configurations" on page 130 and "Changing appliance names in disaster recovery configurations" on page 146.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type `system` to enter system settings mode.
4. Type the following command to specify the appliance name:

   `name` *`identifier`*

   Where *identifier* is a string up to 127 characters long. Use letters and numbers in the name, avoid special characters. The name should not consist of numbers only.
5. After you configure the name, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Querying the appliance name by using the command line

You can discover the name of the appliance by using the command interface.

About this task

To query the appliance name from the command line, you enter system settings mode and enter the required **`show`** command.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode. (You can also run this command from the mqa# prompt.)
3. Type the following command to query the appliance name:

   `show system`

   The appliance responds by displaying system information, which includes the appliance name. For example:

```
                description: IBM MQ Appliance
              serial number: 0000000
             entitlement id: 0000000
                 product id: 5725 [Rev None]
                        OID: 1.3.6.1.4.1.14685.1.8
                     uptime: 0 days 18:01:26
                    contact: (unknown)
                       name: MPW1
                   location: (unknown)
                   services: 72
                 backup mode: normal
                product mode: normal
               login-message:
             custom-ui-file: store:///dp-user-interface-demo.xml
              audit-reserve: 40 kBytes
                     locale: en
system-log-fixed-format: off
```

Configuring appliance users

Appliance users can log in locally to the IBM MQ Appliance web UI and modify all appliance settings, including IBM MQ configuration and objects. Appliance users are created by using the appliance user commands or the Appliance section of the IBM MQ Appliance web UI.

You can also create user groups and assign users to those groups. You can specify privileges for different groups that define what members of those groups can do.

Configuring appliance users by using the IBM MQ Appliance web UI

You can configure appliance users by using the IBM MQ Appliance web UI.

About this task

You use the Manage Appliance part of the IBM MQ Appliance web UI to configure appliance users on your IBM MQ Appliance. The IBM MQ Appliance web UI itself contains detailed help on the fields that you need to configure. Here we describe the major steps that you need to complete.

Note: If you use messaging users and groups on the appliance for authentication records in an HA queue manager, you must set up the same messaging users and groups on all appliances in the HA group. The users and groups are not automatically replicated between the appliances.

Procedure

1. Start the IBM MQ Appliance web UI and click **Manage Appliance**.
2. Select **Manage** > **User access** and then click **User Account**.
3. Click **New** to add a user.
4. Specify the name of the user that you are configuring. The name can contain a maximum of 128 characters. The following characters are valid:
   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name that consists of a single period, or including two periods together, is not permitted)
5. Optionally enter a comment.
6. Specify a password for the user.
7. Set the user level of the account, choose from:
   - **Privileged**. This level is the highest level of access.
   - **Group defined**. The user takes their access level from the group they belong to.
8. Optionally add SNMPv3 credentials.
9. Click **Apply** to save the new user configuration.

Configuring appliance users by using the command line

You can configure appliance users by using the command interface.

About this task

To configure an appliance user from the command line, you enter user configuration mode, specifying the name of the user that you want to configure, and enter the required user configuration commands.

Note: If you use messaging users and groups on the appliance for authentication records in an HA queue manager, you must set up the same messaging users and groups on all appliances in the HA group. The users and groups are not automatically replicated between the appliances.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure an appliance user:

   `user` *name*

   Where *name* identifies the user that you want to configure. If you are creating a new user, *name* can contain up to 128 characters. The following characters are valid:
   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name that consists of a single period, or including two periods together, is not permitted)
4. Use the following user configuration commands to configure the user:

Command	Description
"**admin-state**" on page 437	Sets the administrative state for the configuration.
"**access-level**" on page 450	Assigns an account type to an account.
"**group**" on page 450	Associates an account with a user group.
"**password**" on page 451	Manages the password for an account.

5. After you configure the user, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Configuring appliance user groups by using the IBM MQ Appliance web UI

You can configure appliance user groups by using the IBM MQ Appliance web UI.

About this task

You use the Manage Appliance part of the IBM MQ Appliance web UI to configure appliance user groups on your IBM MQ Appliance. The IBM MQ Appliance web UI itself contains detailed help on the fields that you need to configure. Here we describe the major steps that you need to complete.

Note: If you use messaging users and groups on the appliance for authentication records in an HA queue manager, you must set up the same messaging users and

groups on all appliances in the HA group. The users and groups are not automatically replicated between the appliances.

Procedure

1. Start the IBM MQ Appliance web UI and click **Manage Appliance**.
2. Select **Manage** > **User access** and then click **User Group**.
3. Click **New** to add a user group.
4. Specify the name of the user group that you are configuring. The name can contain a maximum of 128 characters. The following characters are valid:
   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name that comprises a single period, or including two periods together, is not permitted)
5. Optionally enter a comment.
6. Construct an access profile for the group. Enter the set of access policies that define the access privileges for members of the group.
7. Optionally specify which CLI command groups that members of the group can use.
8. Click **Apply** to save the new user configuration.

Example

The following example shows how to provide CLI access to a user group and allow only SSH login for users in that group.

1. Add the following Access profile to block access to all the resources:
   `*/*/*`
2. Add the following Access profile to provide read only access for SSH service:
   `*/*/login/ssh?Access=r`
3. Provide `mqcli` command line access to the user group by adding 'Configuration' in **CLI command groups**.

Configuring appliance user groups by using the command line

You can configure appliance user groups by using the command line interface.

About this task

To configure an appliance user group from the command line, you enter user group configuration mode, specifying the name of the user group that you want to configure, and enter the required user group configuration commands.

Note: If you use messaging users and groups on the appliance for authentication records in an HA queue manager, you must set up the same messaging users and groups on all appliances in the HA group. The users and groups are not automatically replicated between the appliances.

Procedure

1. Connect to the IBM MQ Appliance as described in “Command line access” on page 97. Log in as an administrative user.

2. Type `config` to enter global configuration mode.
3. Type the following command to configure an appliance user:
   `usergroup` *name*

   Where *name* identifies the user group that you want to configure. If you are creating a new user group, *name* can contain up to 128 characters. The following characters are valid:
   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name that comprises a single period, or including two periods together, is not permitted)
4. Use the following user configuration commands to configure the user group:

Command	Description
"**access-policy**" on page 451	Assigns an access policy to the group.
"**add**" on page 453	Adds a command suite to the group capabilities.
"**admin-state**" on page 437	Sets the administrative state for the configuration.
"**delete**" on page 453	Removes a command suite from the group capabilities.

5. After you configure the user, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Configuring queue managers

You can use a number of methods to configure queue managers on the IBM MQ Appliance.

Using the IBM MQ Console

You can use the graphical interface that is provided by the IBM MQ Console to create and configure queue managers and associated objects. For details on how to use the IBM MQ Console, see "Using the IBM MQ Console" on page 157.

Using the appliance command line interface

The appliance provides a number of CLI commands for directly configuring queue managers. For details of these commands, see "Queue manager commands" on page 389.

Using the appliance command line interface to work with the `qm.ini` file

The appliance also provides CLI commands that you can use to edit the `qm.ini` file. You are most likely to edit the `qm.ini` file when you move existing queue managers from other platforms to the appliance as part of consolidating your IBM MQ estate. See "Editing qm.ini files" on page 245.

Configuring environment variables

You can configure environment variables either globally or for a specific queue manager.

Adding an environment variable

You can add an environment variable by using the **setmqvar** command on the command line. You can add either a global environment variable, or a queue manager specific environment variable.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Add the environment variable:
   - To add a global environment variable, enter the following command:

     ```
     setmqvar -k Name -v Value
     ```

     Where:

     Name Specifies the name of the global environment variable.

     Ensure that the value of *Name* is correct before you use the command to add an environment variable. The value of *Name* is not validated.

     Value Specifies the value of the specified environment variable.

     If *Value* is a string that contains spaces, it must be enclosed in double quotation marks. Any double quotation marks that are used in the *Value* must be escaped by using a backslash (\ ).

     Ensure that the value of *Value* is correct before you use the command to add an environment variable. The value of *Value* is not validated.
   - To add an environment variable for a specific queue manager, enter the following command:

     ```
     setmqvar -m QMgrName -k Name -v Value
     ```

     Where:

     QMgrName
     Specifies the queue manager for which the environment variable is added.

     Name Specifies the name of the queue manager environment variable.

     Ensure that the value of *Name* is correct before you use the command to add an environment variable. The value of *Name* is not validated.

     Value Specifies the value of the specified environment variable.

     If *Value* is a string that contains spaces, it must be enclosed in double quotation marks. Any double quotation marks that are used in the *Value* must be escaped by using a backslash (\ ).

     Ensure that the value of *Value* is correct before you use the command to add an environment variable. The value of *Value* is not validated.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the addition of the global environment variable MQSSLRESET with a value of 0:

```
setmqvar -k MQSSLRESET -v 0
```

Modifying an environment variable

You can modify an environment variable by using the **setmqvar** command on the command line. You can modify either a global environment variable, or a queue manager specific environment variable.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. Modify the environment variable:
   - To modify a global environment variable, enter the following command:
     ```
     setmqvar -k Name -v Value
     ```
     Where:

     Name Specifies the name of the global environment variable to modify.

     Ensure that the value of *Name* is correct before you use the command to modify an environment variable. The value of *Name* is not validated.

     Value Specifies the value of the specified environment variable.

     If *Value* is a string that contains spaces, it must be enclosed in double quotation marks. Any double quotation marks that are used in the *Value* must be escaped by using a backslash (\ ).

     Ensure that the value of *Value* is correct before you use the command to modify an environment variable. The value of *Value* is not validated.
   - To modify an environment variable for a specific queue manager, enter the following command:
     ```
     setmqvar -m QMgrName -k Name -v          Value
     ```
     Where:

     QMgrName
     Specifies the queue manager for which the environment variable is modified.

     Name Specifies the name of the queue manager environment variable.

     Ensure that the value of *Name* is correct before you use the command to modify an environment variable. The value of *Name* is not validated.

     Value Specifies the value of the specified environment variable.

     If *Value* is a string that contains spaces, it must be enclosed in double quotation marks. Any double quotation marks that are used in the *Value* must be escaped by using a backslash (\ ).

     Ensure that the value of *Value* is correct before you use the command to modify an environment variable. The value of *Value* is not validated.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the modification of the global environment variable `MQSSLRESET` with a value of `1000`:

```
setmqvar -k MQSSLRESET -v 1000
```

Removing an environment variable

You can remove an environment variable by using the **setmqvar** command on the command line. You can remove a global environment variable, or a queue manager specific environment variable.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Remove the environment variable:
   - To remove a global environment variable, enter the following command:

     ```
     setmqvar -k Name -d
     ```

     Where:

     Name Specifies the name of the global environment variable.

     Ensure that the value of *Name* is correct before you use the command to remove an environment variable. The value of *Name* is not validated.
   - To remove an environment variable from a specific queue manager, enter the following command:

     ```
     setmqvar -m QMgrName -k Name -d
     ```

     Where:

     QMgrName
     Specifies the queue manager for which the environment variable is removed.

     Name Specifies the name of the queue manager environment variable.

     Ensure that the value of *Name* is correct before you use the command to remove an environment variable. The value of *Name* is not validated.
3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the removal of the global environment variable `MQSSLRESET`:

```
setmqvar -k MQSSLRESET -d
```

Viewing environment variables

You can view environment variables by using the **dspmqvar** command on the command line. You can view either global environment variables, or queue manager specific environment variables.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. View one or more environment variables:
   - To view all global environment variables, enter the following command:

     ```
     dspmqvar
     ```

   - To view a specific global environment variable, enter the following command:

     ```
     dspmqvar -k Name
     ```

     Where:

     Name Specifies the name of the global environment variable to view.

     Ensure that the value of *Name* is correct before you use the command to view an environment variable. The value of *Name* is not validated.

   - To view all environment variables for a specific queue manager, enter the following command:

     ```
     dspmqvar -m QMgrName
     ```

     Where:

     QMgrName
     Specifies the queue manager for which the environment variable is viewed.

   - To view a specific environment variable for a specific queue manager, enter the following command:

     ```
     dspmqvar -m QMgrName -k Name
     ```

     Where:

     QMgrName
     Specifies the queue manager for which the environment variable is viewed.

     Name Specifies the name of the queue manager environment variable to view.

     Ensure that the value of *Name* is correct before you use the command to view an environment variable. The value of *Name* is not validated.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example views of the global environment variable `MQSSLRESET`:

```
dspmqvar -k MQSSLRESET
```

Configuring high availability

You can configure a pair of appliances to provide a high availability (HA) solution.

For more information about HA on the IBM MQ Appliance, see “High availability” on page 3.

To configure HA for a pair of appliances, you must complete the following steps:

1. Connect the appliances together. For more information, see “Configuring the hardware for high availability.”
2. Create an HA group for the appliances. For more information, see “Configuring the high availability group” on page 130.
3. Create HA queue managers on the appliances. For more information, see “Configuring high availability queue managers” on page 134.

Messaging users that connect to HA queue managers must be recognized on both appliances in the HA group. You can use an external LDAP directory to store details of messaging users, which can be accessed by both appliances in the HA group. Otherwise, you must set up local messaging users on both appliances.

You connect your applications to HA queue managers by using IBM MQ clients that connect to the appliance. You must define connection details for both the appliances in the HA group. For guidance on how to do this procedure, see Channel and client reconnection in the IBM MQ documentation.

View the video for a demonstration of configuring a pair of IBM MQ Appliances to provide an HA solution:

Configuring an HA Group on Two IBM MQ Appliances

Configuring the hardware for high availability

You must physically connect the appliances together before you can configure the IBM MQ Appliance for high availability (HA).

Before you begin

You need three Ethernet cables of sufficient length to directly connect the two appliances together. Cables are provided with the appliance for this purpose, see “Connect the appliance to a network” on page 53. If you supply your own cables, note that these need to be straight through cables, not crossover cables.

Alternatively, you can connect two appliances that are further apart by using a switch. If you choose to use a switch, ensure that the primary and alternative interfaces are connected using independent switches. Using independent switches increases the reliability of communication within the HA group as there is not a single point of failure.

If neither the HA group primary interface nor the HA group alternate interface are connected, the appliances in the HA group are unable to determine the state of the other appliance. This situation can cause high availability (HA) queue managers to run on both appliances simultaneously.

About this task

The IBM MQ Appliance uses three Ethernet ports on each appliance to configure HA. These ports are shown in the following diagrams:

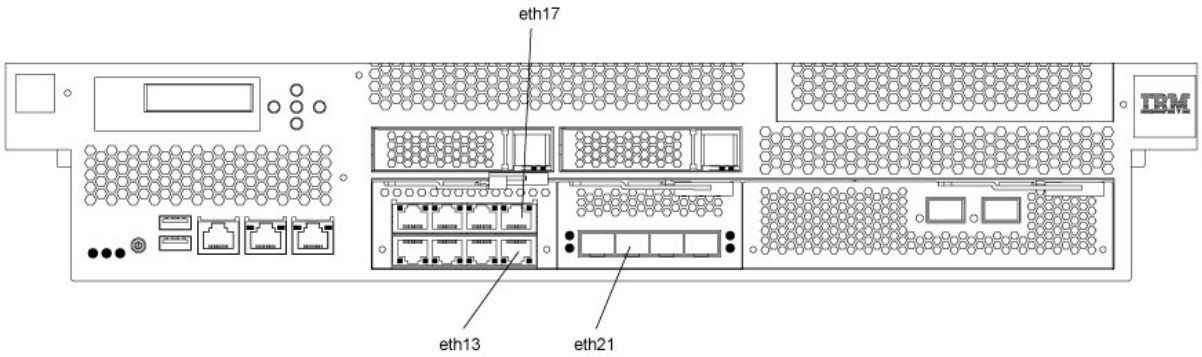

Figure 32. Ethernet ports on M2001 appliance

Note: Do not disable these interfaces through the IBM MQ Appliance web UI or the command line interface in an attempt to test the high availability features. Disabling the links causes unexpected results; you cannot simulate real high availability behavior by disabling a network interface.

Procedure

1. Connect an Ethernet cable between port eth13 on the first appliance and eth13 on the second appliance. This connection is the **HA group primary interface**.
2. Connect an Ethernet cable between port eth17 on the first appliance and eth17 on the second appliance. This connection is the **HA group alternate interface**.
3. Connect an Ethernet cable between port eth21 on the first appliance and eth21 on the second appliance. This connection is the **replication interface**.
4. Ensure that the three Ethernet ports have IP addresses configured. If the ports were not configured when you initialized the appliances, then configure them by using the procedure that is described in "Ethernet interfaces" on page 103. You can use the IBM MQ Appliance web UI or the command line to configure the interfaces.
5. Ensure that both appliances have system names. If the names were not assigned when you initialized the appliances, then name them by using the procedure that is described in "Configuring the appliance name" on page 118. You can use the IBM MQ Appliance web UI or the command line to assign appliance names.

What to do next

After the appliances are connected, you can create an HA group for the appliances. For more information, see "Creating a high availability group" on page 130.

Changing IP addresses in high availability configurations

If you change the IP addresses of any of the Ethernet ports of the appliances in a high availability configuration, high availability operation is no longer available and data is likely to be partitioned.

High availability (HA) configurations use the eth13, eth17, and eth21 ports on both appliances in the HA. If you need to change IP addresses for any of these ports, you must use the following procedure:

1. Remove the HA configuration on both appliances. You remove HA by removing queue managers from HA control, see "Removing a queue manager

from a high availability group" on page 137, and then removing the HA group itself, see "Deleting a high availability group" on page 133.

2. Allocate new IP numbers to the Ethernet ports as required, observing the requirement that the addresses are on the same, dedicated subnet. See "Configuring Ethernet interfaces by using the command line" on page 105.
3. Recreate the HA configuration by configuring the queue managers as described in "Creating a high availability group" and "Creating a high availability queue manager" on page 134.

Changing appliance names in high availability configurations

If you change the appliance name of either or both of the appliances in a high availability configuration, high availability operation is no longer available and data is likely to be partitioned.

If you need to change appliance names, you must use the following procedure:

1. Remove the HA configuration on both appliances. You remove HA by removing queue managers from HA control, see "Removing a queue manager from a high availability group" on page 137, and then removing the HA group itself, see "Deleting a high availability group" on page 133.
2. Allocate new appliance names as required, see "Configuring the appliance name by using the command line" on page 118.
3. Recreate the HA configuration by configuring the queue managers as described in "Creating a high availability group" and "Creating a high availability queue manager" on page 134.

Configuring the high availability group

When you configure appliances to be part of a high availability (HA) solution, you must configure the appliances to be part of the HA group.

The HA group controls the availability of the queue managers within the group, determining where the queue managers run. You must create the HA group before you can create HA queue managers to run in the group. After you create the HA group, you can view the status of the group. You can also view which queue managers are in the HA group.

If you use messaging users and groups on the appliance for authentication records in an HA queue manager, you must set up the same messaging users and groups on both appliances. The users and groups are not automatically replicated between the appliances.

Creating a high availability group

You can create a high availability (HA) group by using the **crthagrp** command on the command line.

Before you begin

Before you can create an HA group, you must configure the appliances that you want to group. For more information, see "Configuring the hardware for high availability" on page 128

About this task

The HA group controls the availability of the queue managers within the group, determining where the queue managers run. You must create the HA group before

you can create HA queue managers to run in the group. After you create the HA group, you can view the status of the group. You can also view which queue managers are in the HA group.

To create an HA group and generate unique keys for communication between the appliances in the group, you must enter commands on both appliances in the group.

Note: The messaging users and groups in authentication records in an HA queue manager must be available on all appliances in the HA group. Because the users and groups in the internal user store are not automatically replicated between the appliances, it is recommended that you use an external LDAP repository for HA queue managers.

Procedure

1. Enter the IBM MQ administration mode on both appliances by entering the following command:

   ```
   mqcli
   ```

2. On the first appliance (machine A), enter the following command:

   ```
   prepareha -s SecretText -a IPAddressOfMachineB [-t timeout]
   ```

   where:

   -s *SecretText*
   Specifies a string that is used to generate a short-lived password. The password is used to set up the unique key for the two appliances.

   -a *IPAddressOfMachineB*
   Specifies the IP address of the HA primary interface on the second appliance in the group. You must specify the IP address using ip v4 dotted decimal notation (for example, “192.0.2.8”).

   -t *timeout*
   Optionally specifies a timeout in seconds. The appliance waits silently for this period for the second appliance to contact it. If you do not specify a timeout, the appliance waits for ten minutes.

3. On the second appliance (machine B), enter the following command:

   ```
   crthagrp -s SecretText -a IPAddressOfMachineA
   ```

   where:

   -s *SecretText*
   Specifies the same string that was specified in the **prepareha** command on machine A.

   -a *IPAddressOfMachineA*
   Specifies the IP address of the HA primary interface on the first appliance in the group. You must specify the IP address using ip v4 dotted decimal notation (for example, “192.0.2.8”).

Example

The following example shows the creation of an HA group for appliances `appl1` and `appl2` where a new, unique key is generated for communication between the appliances. The HA group primary interface of `appl2` has the IP address `192.0.2.8`; the HA group primary interface of `appl1` has the IP address `192.0.2.7`.

The following command is run from `appl1`:

```
prepareha -s AGEW1823510HH -a 192.0.2.8
```

The following command is run from `appl2`:

```
crthagrp -s AGEW1823510HH -a 192.0.2.7
```

What to do next

After the HA group is created, you can create HA queue managers that are controlled by the HA group. For more information, see "Creating a high availability queue manager" on page 134.

Viewing the status of appliances in a high availability group

You can view the status of appliances in a high availability (HA) group by using the command line or the IBM MQ Console.

Viewing the high availability status by using the command line:

You can view the status of appliances in a high availability (HA) group by using the **dsphagrp** command on the command line.

About this task

The **dsphagrp** command returns information about the operational status of each of the appliances in the HA group. The status can be one of the following statuses:

- `Online`. The appliance is available.
- `Offline`. The appliance is unavailable.
- `Standby`. The appliance has been temporarily removed from the HA group.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. View the status of the appliances in the HA group by entering the following command from one of the appliances:
   ```
   dsphagrp
   ```
3. Optional: Exit the IBM MQ administration mode by entering the following command:
   ```
   exit
   ```

Viewing the high availability status by using the IBM MQ Console:

You can view the status of appliances in a high availability (HA) group by using the IBM MQ Console.

About this task

The appliance can have one of the following statuses:

- `Online`. The appliance is available.
- `Offline`. The appliance is unavailable.
- `Standby`. The appliance has been temporarily removed from the HA group.

Procedure

1. Start the IBM MQ Appliance web UI and click **MQ Console**.

2. Click the High Availability menu in the console title bar. The menu displays the status of both appliances in the group.

Deleting a high availability group

You can delete an existing high availability (HA) group by using the command line or the IBM MQ Console.

Deleting a high availability group by using the command line:

You can delete an existing high availability (HA) group by using the **dlthagrp** command on the command line.

About this task

You must delete all HA queue managers in the HA group before you delete the group.

You run the command on one appliance and the HA group is deleted on both appliances in the group. If the other appliance is not available at the time of the delete, the command must be entered on the other appliance to delete the group on that appliance.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. Enter the following command to delete the HA group:
   ```
   dlthagrp
   ```

Deleting a high availability group by using the IBM MQ Console:

You can delete an existing high availability (HA) group by using the IBM MQ Console.

About this task

You must delete all HA queue managers in the HA group before you delete the group.

You use the console on one appliance and the HA group is deleted on both appliances in the group. If the other appliance is not available at the time of the delete, the group must be deleted on that appliance when it is next available.

Procedure

1. Start the IBM MQ Appliance web UI and click **MQ Console**.

2. Click the High Availability menu in the console title bar, and select **Delete group**.

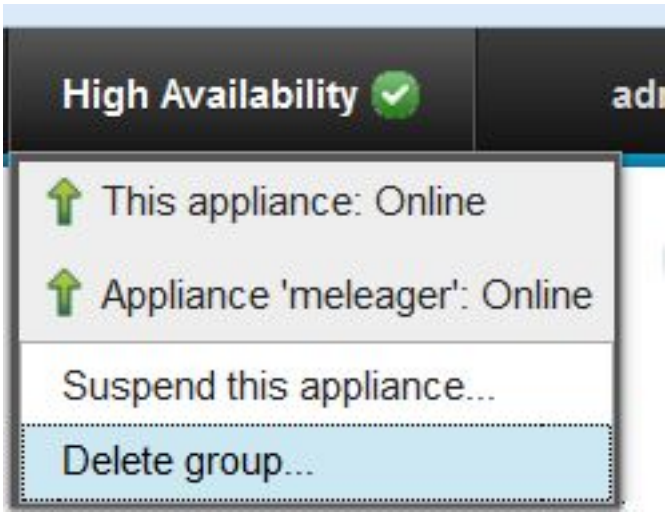

3. A window prompts you to confirm the deletion. Click **Delete**.

Configuring high availability queue managers

When you configure queue managers, you can specify that they belong to a high availability group.

Creating a high availability queue manager

You can create a high availability (HA) queue manager by using the **crtmqm** command on the command line. After the queue manager is created, it is automatically started under the control of the HA group.

Before you begin

Before you can create an HA queue manager on an appliance, you must add the appliance to an HA group. For more information, see "Creating a high availability group" on page 130.

About this task

You create a queue manager and specify that it is part of an HA group. Each HA queue manager that you create uses a unique port on the HA replication interface. The first HA queue manager created uses port 7789, the second 7790, and so on.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Create the HA queue manager by using one of the following commands:
   - Create the HA queue manager with the default file system size of 64 GB by entering the following command:

     ```
     crtmqm -sx QMgrName
     ```

     Where:

     QMgrName
     Specifies the name of the HA queue manager that you want to create.

- Create the HA queue manager with a specific file system size by entering the following command:

  ```
  crtmqm -sx -fs FileSystemSize          QMgrName
  ```

  Where:

  FileSystemSize
  Specifies the file system size that the queue manager is created with.

  This value is a numeric value, which is specified in GB.

  QMgrName
  Specifies the name of the HA queue manager that you want to create.

Note:

- The HA queue manager is created on the appliance on which the **crtmqm** command is run. The queue manager automatically starts on that appliance after it is created. You cannot use the **strmqm** command to start the queue manager.
- You can use other **crtmqm** parameters in the command. For more information about the available parameters, see "crtmqm" on page 326.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the creation of an HA queue manager HAQM1:

```
crtmqm -sx HAQM1
```

Adding an existing queue manager to a high availability group

You can add an existing queue manager to a high availability (HA) group by using the **sethagrp** command on the command line.

Before you begin

Before you can add an existing queue manager to a group, the queue manager must be stopped.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Enter the following command to stop the queue manager:

   ```
   endmqm QMname
   ```

3. Enter the following command to add an existing queue manager to the HA group:

   ```
   sethagrp -i QMname
   ```

   Where *QMname* is the name of the existing queue manager. You must check that a queue manager with that name does not already exist on the other appliance. The queue manager is added to the group and is started.

4. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the existing queue manager `QM1` being added to the HA group:

```
sethagrp -i QM1
```

Viewing the status of a high availability queue manager

You can view the status of a queue manager in a high availability (HA) group by using the **status** command on the command line , or by using the IBM MQ Console.

About this task

The **status** command returns information about the operational status of a specified queue manager in the HA group. The status can include the following information:

- The high availability role of the queue manager (reported as `Primary` or `Secondary`).
- The current high availability status:

  Normal
  The appliances in the disaster recovery configuration are operating normally.

  This appliance in standby mode
  This status means that the appliance has been suspended (by using the **sethagrp -s** command).

  Secondary appliance in standby mode
  This status means that the other appliance in the HA pair has been suspended (by using the **sethagrp -s** command).

  Both appliances in standby mode
  This status means that both appliances in the HA pair have been suspended (by using the **sethagrp -s** command).

  Secondary appliance unavailable
  This status means that the connections to the other appliance in the HA pair have been lost.

  Partitioned
  If the two appliances lose the replication interface, the queue managers are immediately flagged as partitioned. The running queue manager might accumulate out-of-sync data. The other queue manager remains in standby with no out-of-sync data. When the connection is remade, the partitioned state is automatically resolved and replication is resumed.

  Synchronization in progress
  This status is displayed when the primary queue manager is replicating data to the secondary queue manager.

  Inactive
  The queue manager is inactive on both appliances in the HA pair.
- The preferred appliance setting for the queue manager, set to `This Appliance` or `Other Appliance`.
- The percentage complete of a synchronization operation. This information is shown only when the status is `Synchronization in progress`.
- The estimated time at which a synchronization will complete. This information is shown only when the status is `Synchronization in progress`.

- The amount of out-of-sync data that exists on this instance of the queue manager. This is the amount of data written to this instance of the queue manager since it entered the partitioned state. This information is shown only when the status is `Partitioned`.

Procedure

- To view the HA status of a queue manager by using the command line interface:
  1. Enter the IBM MQ administration mode by entering the following command:
     ```
     mqcli
     ```
  2. View the status of an HA queue manager by entering the following command from one of the appliances:
     ```
     status QMgrName
     ```
     Where:

     QMgrName
     Specifies the name of the HA queue manager that you want to view the status of.
  3. Exit the IBM MQ administration mode by entering the following command:
     ```
     exit
     ```
- To view the HA status of a queue manager by using the IBM MQ Console:
  1. Open the console and find the widget that displays the queue manager.
  2. Select the queue manager in the widget and select the properties icon from the tool bar .
  3. In the properties window, click on the **High availability status** section to open it.

Removing a queue manager from a high availability group

You can remove a queue manager from a high availability (HA) group and run it as a standalone queue manager by using the **sethagrp** command on the command line.

Before you begin

If the queue manager is also part of a disaster recovery (DR) configuration, you must remove it from the DR configuration before you remove it from the HA group. See “Removing a queue manager from a disaster recovery configuration by using the command line” on page 149.

About this task

You must run the command on the queue manager primary appliance (the appliance that the queue manager is running on). You can discover where the queue manager is running by using the **dspmq** command or the **status** *qmanager* command. Either command will report the status as `Running` for the current appliance, or `Running elsewhere` for the other appliance in the HA group.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. Enter the following command to stop the queue manager:
   ```
   endmqm QMname
   ```

3. Enter the following command to remove the queue manager from the HA group and run it as a standalone queue manager:

   ```
   sethagrp -e QMname
   ```

   Where *QMname* is the name of the queue manager. The queue manager is removed from the group. You must then use the **strmqm** command to start the queue manager.
4. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the queue manager HAQM1 being removed from the HA group:

```
sethagrp -e HAQM1
```

Example network set up for HA configuration

The example shows the network configuration for an high availability implementation.

The configuration is illustrated in the following diagram. The two HA appliances are located in adjacent racks, and are directly connected to each other with the supplied cables. It is recommended that the IP addresses are in separate subnets for each connection (as shown in the diagram). Otherwise you need to take other steps to ensure that data leaves the appliance on the correct Ethernet interface (for example, by setting up routes on the appliance).

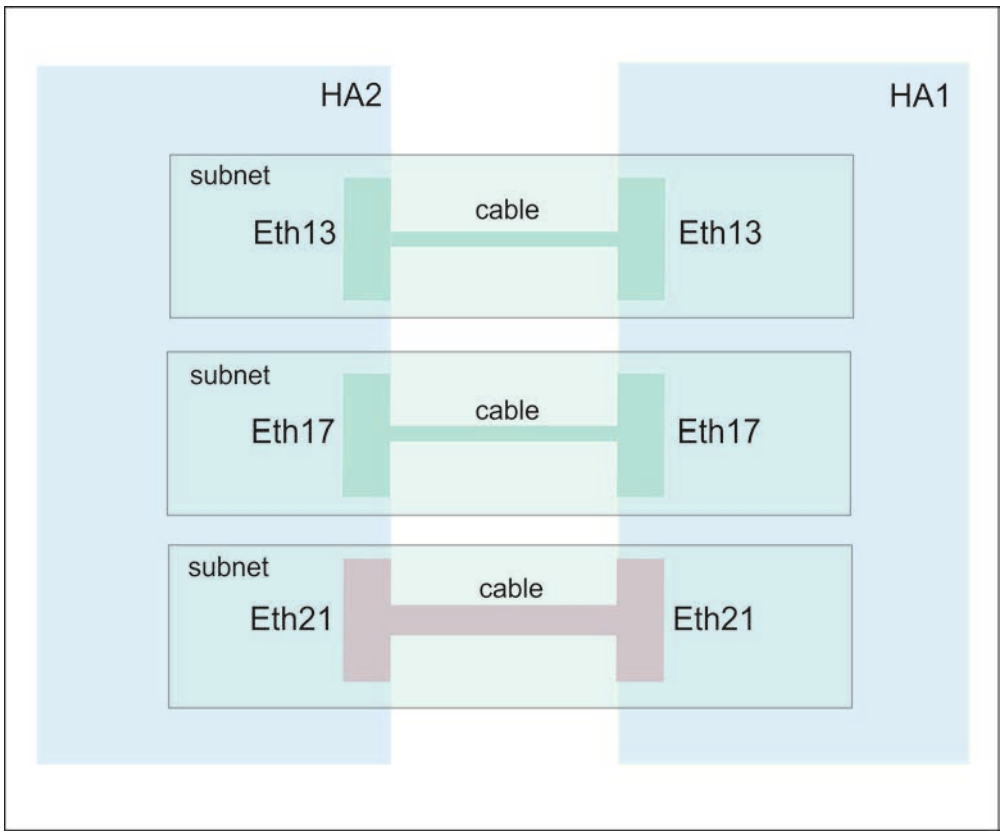

Figure 33. Example HA/DR network configuration

The following tables show how the Ethernet ports are configured on the two HA appliances, HA1 and HA2.

Table 13. Ethernet ports on appliance HA1

Ethernet port	Example IP address	Port	Description
eth13	192.0.10.11	5404, 5405 for heartbeat, 2222 for configuration	HA primary group interface
eth17	192.0.20.11	5404, 5405 for heartbeat, 2222 for configuration	HA group alternative interface
eth21	192.0.30.11	Each HA queue manager uses a port, starting at 7789 for the first created, 7790 for the second created, and so on, up to 8021	HA replication interface
eth22	203.0.113.0		Data interface

Table 14. Ethernet ports on appliance HA2

Ethernet port	Example IP address	Port	Description
eth13	192.0.10.12	5404, 5405 for heartbeat, 2222 for configuration	HA primary group interface

Table 14. Ethernet ports on appliance HA2 (continued)

Ethernet port	Example IP address	Port	Description
eth17	192.0.20.12	5404, 5405 for heartbeat, 2222 for configuration	HA group alternative interface
eth21	192.0.30.12	Each HA queue manager uses a port, starting at 7789 for the first created, 7790 for the second created, and so on, up to 8021	HA replication interface
eth22	203.0.113.2		Data interface

Configuring disaster recovery for a high availability queue manager

You can specify that a high availability queue manager also belongs to a disaster recovery configuration.

Both appliances in a high availability pair are typically located in the same data center. If some disaster befalls the data center, and both appliances are unavailable, you can manually start the queue manager on a recovery appliance located in another data center. See “Disaster recovery for a high availability configuration” on page 6 for an overview.

To configure disaster recovery for a high availability queue manager, complete the following steps:

1. On the appliance where your HA queue manager is running, enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Stop the HA queue manager:

   ```
   endmqm queue_manager
   ```

3. Specify that the queue manager is the primary instance in a disaster recovery configuration and include a floating IP address that can be used by either of the appliances in the HA pair:

   ```
   crtdrprimary –m queue_manager –r RecoveryName –i RecoveryIP
   –p port_number -f floating_IP
   ```

   Where:

   -m ***QMName***
   Specifies the queue manager that you are preparing for participation in a disaster recovery configuration.

   -r ***RecoveryName***
   Specifies the name of the IBM MQ Appliance that is the recovery appliance.

   -i ***RecoveryIP***
   Specifies the IP address of the recovery appliance.

-p *port*

Specifies the port that the data replication listener on each appliance uses.

-f *floatingIP*

The floating IP address is an IPv4 address that is used to replicate queue manager data from whichever HA appliance the queue manager is currently running on to the queue manager on the recovery appliance. The floating IP address must be in the same subnet group as the static IP address assigned to the replication port (eth20) on both appliances.

Note that you do not physically configure an Ethernet port with this address. Select a free IP address in the same subnet as the replication ports on the two appliances, and specify it in the **crtdrprimary** command to make it the IP used for replication with the recovery appliance. You must specify a different floating IP address for each of the HA queue managers that you configure disaster recovery for.

The `crtdrprimary` command configures the queue manager on both appliances in the HA pair, and reserves storage for the data snapshot on both appliances.The **crtdrprimary** command returns a **crtdrsecondary** command when it has completed, for example:

```
Queue manager QM3 is prepared for Disaster Recovery replication.
Now execute the following command on appliance mydrappl:
crtdrsecondary –m QM3 –s 65536 –l myliveapp3 –i 198.51.100.10 –p 2015
```

4. Copy the **crtdrsecondary** command and run it on the recovery appliance. This creates a secondary version of the queue manager, and queue manager data is replicated from the primary queue manager.

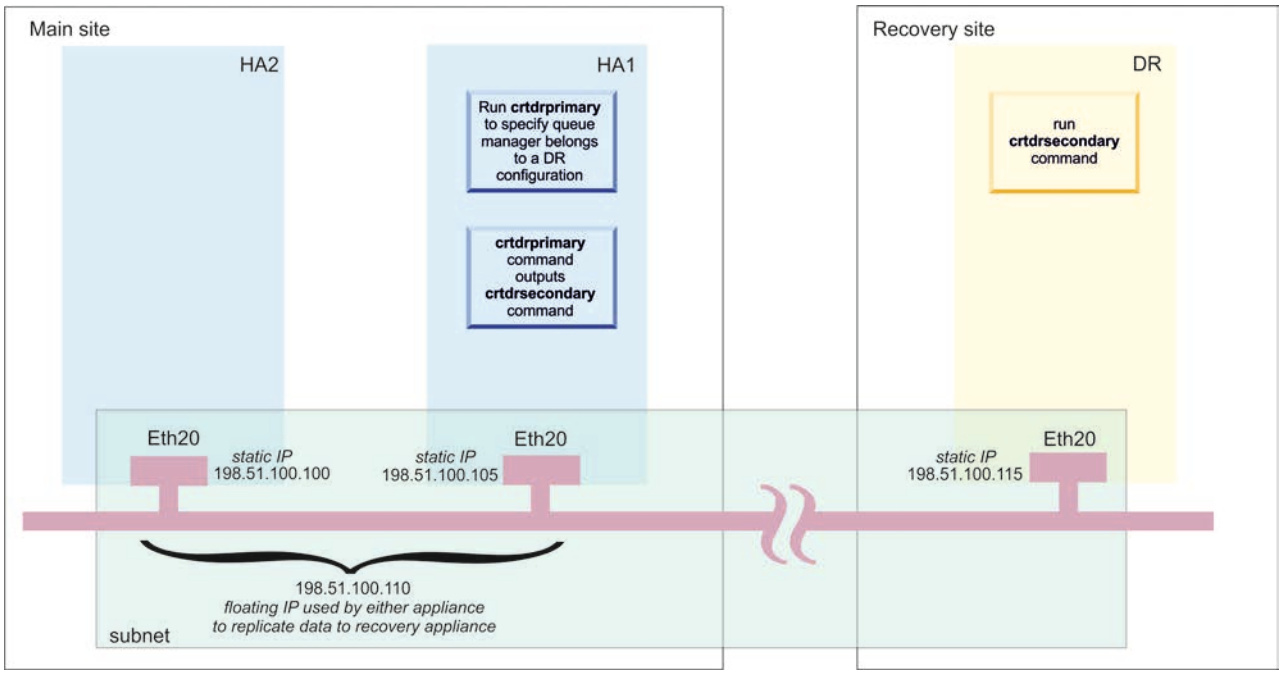

Figure 34. Configuring an HA queue manager for disaster recovery (using example IP addresses)

Example network set up for HA/DR configuration

The example shows the network configuration for an HA pair at the main site with a DR appliance at a recovery site.

The configuration is illustrated in the following diagram. The two HA appliances are located in adjacent racks, and are directly connected to each other with the supplied cables. The replication connection to the recovery appliance and the data connection for applications are made by way of a switch.

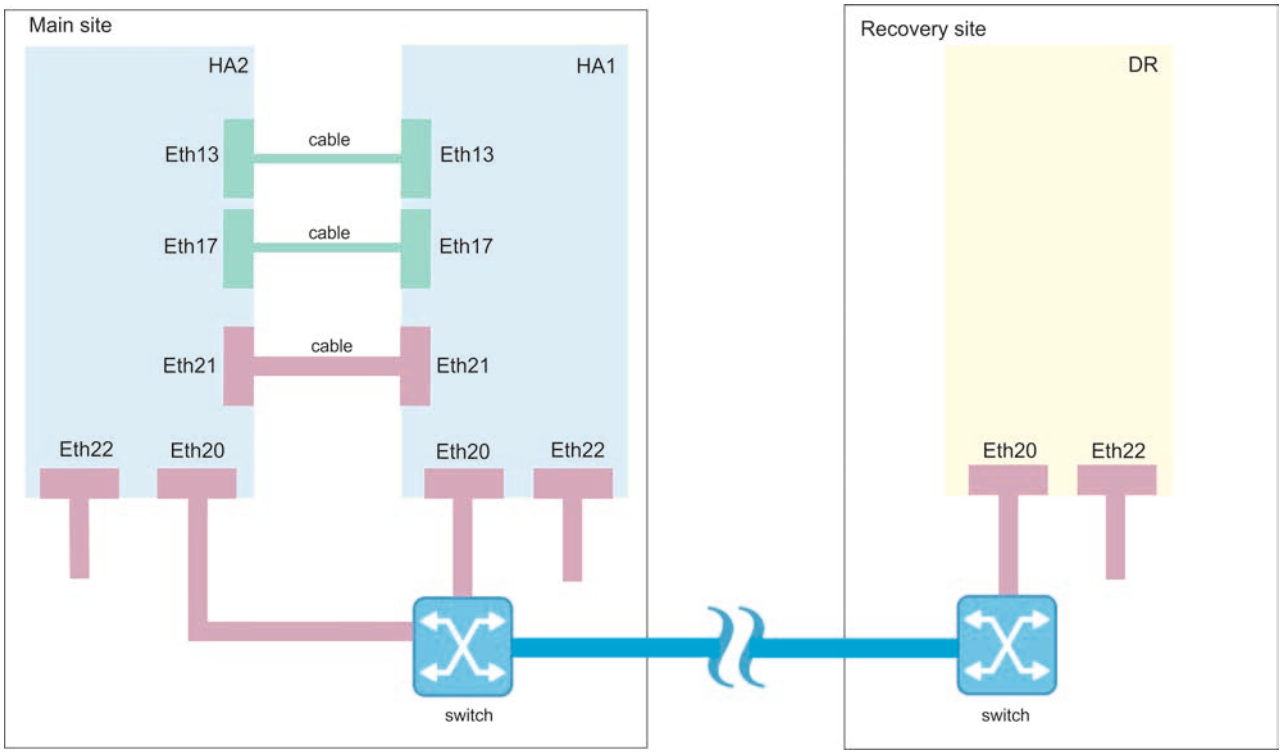

Figure 35. Example HA/DR network configuration

The following tables show how the Ethernet ports are configured on the two HA appliances, HA1 and HA2, and the DR appliance, DR.

Table 15. Ethernet ports on appliance HA1

Ethernet port	IP address	Description
eth13	192.0.10.11	HA primary group interface
eth17	192.0.20.11	HA group alternative interface
eth20	192.0.40.11	Static IP address configured for DR port
-	192.0.40.19	Floating IP address used by appliance HA1 or HA2 for replication with DR appliance
eth21	192.0.30.11	HA replication interface
eth22	203.0.113.0	Data interface

Table 16. Ethernet ports on appliance HA2

Ethernet port	IP address	Description
eth13	192.0.10.12	HA primary group interface
eth17	192.0.20.12	HA group alternative interface
eth20	192.0.40.12	Static IP address configured for DR port
-	192.0.40.19	Floating IP address used by appliance HA1 or HA2 for replication with DR appliance
eth21	192.0.30.12	HA replication interface
eth22	203.0.113.2	Data interface

Table 17. Ethernet ports on appliance DR

Ethernet port	IP address	Description
eth20	192.0.40.13	DR replication interface
eth22	192.0.2.0	Data interface

Replacing a failed node in a high availability group

If an appliance that belongs to a high availability (HA) group fails, you can replace the appliance and then restore the HA group by following this procedure.

Before you begin

When a node in an HA group fails, the queue managers fail over to the remaining appliance in the group. To restore high availability function after you replace or repair the failed appliance, you must first deconstruct the HA group by running the queue managers standalone and deleting the HA group from the remaining appliance. You then create a new HA group, and add the queue managers back to it.

Before you create the new group, you must ensure that both appliances are running the same level of firmware. If your new appliance is running a later version of the firmware, you must either upgrade your existing appliance, or downgrade your new appliance.

Procedure

1. On the appliance that did not fail, stop each queue manager by using the following command:

   ```
   endmqm QMname
   ```

2. If the queue manager is part of a disaster recovery configuration as well as part of an HA group, you must remove it from the disaster recovery configuration. Use the following command:

   ```
   dltdrprimary -m QMname
   ```

3. Enter the following command to remove a queue manager from the HA group and run it as a stand-alone queue manager. The queue manager must be stopped before you run this command.

   ```
   sethagrp -e QMname
   ```

Where *QMname* is the name of the queue manager. The queue manager is removed from the HA group. You can use the **strmqm** command to restart the queue manager and run it in a stand alone configuration while you replace the failed node, if required.

Repeat this command for all HA queue managers.

4. Delete the HA group by entering the following command:

   ```
   dlthagrp
   ```

5. On both the existing appliance and the replacement appliance, create a new HA group by using the **prepareha** and **crthagrp** commands, as described in "Creating a high availability group" on page 130.
6. On the appliance that did not fail, enter the following command to add a queue manager back to the HA group. The queue manager must be stopped before you run this command.

   ```
   sethagrp -i QMname
   ```

   Where *QMname* is the name of the existing queue manager. The queue manager is added to the group and is started. Repeat for all the queue managers that were previously part of the HA group.
7. Set the preferred appliance for the queue manager by running the following command:

   ```
   sethapreferred QMname
   ```

   Repeat this command for each queue manager. Run the command on the appliance that did not fail if you want that appliance to be the preferred location. Run the command on the replaced or repaired appliance if you want that appliance to be the preferred location.
8. If you want to restore disaster recovery capability to any of the queue managers, follow the instructions in "Configuring disaster recovery for a high availability queue manager" on page 140.

Configuring disaster recovery

You can configure a pair of appliances to provide a disaster recovery (DR) solution. You can also configure a disaster recovery solution for a high availability pair, whereby a high availability queue manager can run on a single DR appliance.

For more information about disaster recovery on an IBM MQ Appliance, see "Disaster recovery" on page 6. For more information about disaster recovery for a high availability pair, see "Disaster recovery for a high availability configuration" on page 6.

Configuring the hardware for disaster recovery

You must connect the appliances together via a network before you can configure the IBM MQ Appliance for disaster recovery (DR).

About this task

The IBM MQ Appliance uses a single 10 GB Ethernet port on each appliance to configure DR. The port is shown in the following diagrams:

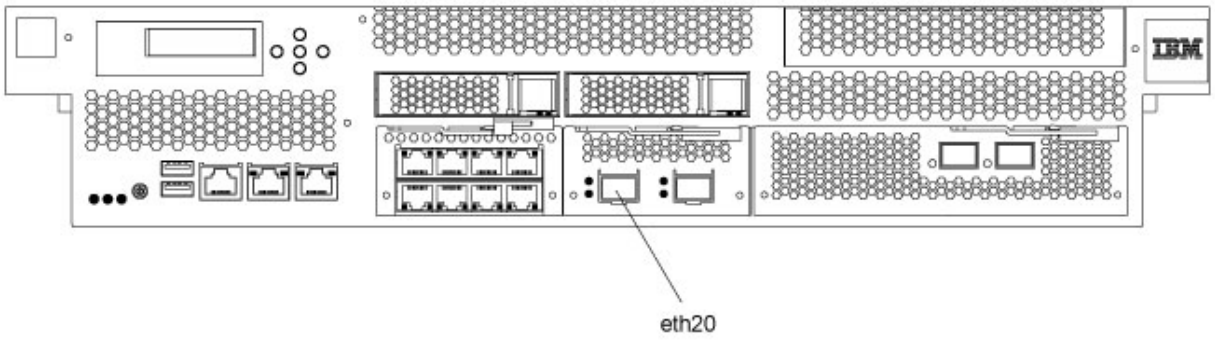

Figure 36. Eth20 on an M2000 appliance

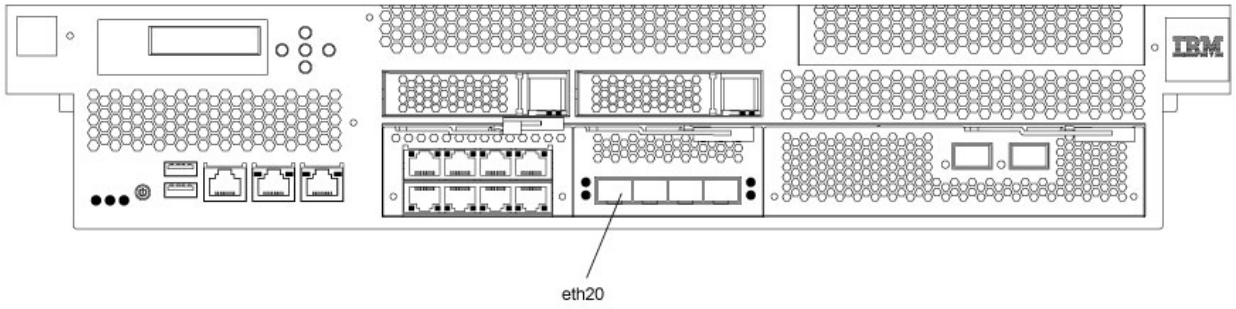

Figure 37. Eth20 on an M2001 appliance

Procedure

1. Connect an Ethernet cable between port eth20 on the appliance and your network. This connection is the **replication interface**.
2. Ensure that the Ethernet port has an IP address configured. If the port was not configured when you initialized the appliance, then configure it by using the procedure that is described in "Ethernet interfaces" on page 103. You can use the IBM MQ Appliance web UI or the command line to configure the interface.
3. Ensure that the Eth20 port connects with the Eth20 port on the other appliance. You can do this either by ensuring that the IP address of each Eth20 port belongs to the same, dedicated subnet, or by defining a static IP route on each appliance between the two Eth20 ports (see "**ip-route**" on page 498).
4. Ensure that both the main and recovery appliances have system names. If the names were not assigned when you initialized the appliances, then name them

by using the procedure that is described in "Configuring the appliance name" on page 118. You can use the IBM MQ Appliance web UI or the command line to assign appliance names.

Changing IP addresses in disaster recovery configurations

If you change the IP addresses of either or both of the eth20 ports of the appliances in a disaster recovery configuration, replication is no longer possible between the two appliances.

If you need to change IP addresses, you must use the following procedure:

1. Remove the disaster recovery (DR) configuration on both appliances. You remove DR by removing both primary and secondary queue managers from DR control, see "Removing a queue manager from a disaster recovery configuration by using the command line" on page 149.
2. Allocate new IP numbers to the two eth20 ports as required, observing the requirement that the addresses are on the same, dedicated subnet. See "Configuring Ethernet interfaces by using the command line" on page 105.
3. Recreate the DR configuration by configuring the queue managers as described in "Configuring queue managers for disaster recovery by using the command line" on page 147.

You can also use the IBM MQ Appliance web UI to perform these operations. Follow the related links for details.

Changing appliance names in disaster recovery configurations

If you change the appliance name of either or both of the appliances in a disaster recovery configuration, replication is no longer possible between the two appliances.

If you need to change appliance names, you must use the following procedure:

1. Remove the disaster recovery (DR) configuration on both appliances. You remove DR by removing both primary and secondary queue managers from DR control, see "Removing a queue manager from a disaster recovery configuration by using the command line" on page 149.
2. Allocate new appliance names as required, "Configuring the appliance name by using the command line" on page 118.
3. Recreate the DR configuration by configuring the queue managers as described in "Configuring queue managers for disaster recovery by using the command line" on page 147.

You can also use the IBM MQ Appliance web UI to perform these operations. Follow the related links for details.

Configuring disaster recovery queue managers

When you configure queue managers, you can specify that they belong to a disaster recovery configuration.

Configuring queue managers for disaster recovery by using the command line

You set up a disaster recovery configuration by setting up a primary and secondary queue manager.

Before you begin

You must have configured a main appliance and a recovery appliance as described in "Configuring the hardware for disaster recovery" on page 144.

About this task

When you configure disaster recovery, you specify that an existing queue manager on your main appliance is the primary queue manager. When that command completes, it outputs a further command that you run on your recovery appliance to create the secondary queue manager. You then run that command on the recovery appliance to create another instance of that queue manager and specify that it is the secondary queue manager.

You must ensure that there is sufficient free memory for the snapshot of the queue manager data that is required for disaster recovery. For example, a queue manager created with the default 64 GB size requires a further 64 GB of free space to be reserved for the snapshot of the queue manager data.

Procedure

1. On the appliance that is designated as your main appliance, run the `crtdrprimary` command, specifying an existing queue manager. The queue manager must be stopped when you run this command.

   ```
   crtdrprimary -m QMName -r RecoveryName -i RecoveryIP -p Port
   ```

   Where:

   -m *QMName*
   : Specifies the queue manager that you are preparing for participation in a disaster recovery configuration. The queue manager must be stopped when you run the command.

   -r *RecoveryName*
   : Specifies the name of the IBM MQ Appliance that is the recovery appliance.

   -i *RecoveryIP*
   : Specifies the IP address of the DR connection (eth20) of the recovery appliance.

   -p *port*
   : Specifies the port that the data replication listener on each appliance uses. The port must be in the range 1025-9999 and must be the same on both appliances (do not use port 2222, it is reserved by the appliance). Each listener is active only on the replication interface (eth20), but you must ensure that the listener does not conflict with any services configured to listen on all appliance interfaces (for example, MQ listeners, or SSH and WebUI services, where these are not restricted to

particular local IP addresses). The data replication listener must also not be blocked by any routing or firewalls between the appliances on the replication network

For example:

```
crtdrprimary –m QM1 –r mydrappl –i 198.51.100.3 –p 2015
```

On successful completion, the command outputs the **crtdrsecondary** command.

You can now restart the queue manager.

2. On the appliance that is designated as your recovery appliance, run the command that was output by the `crtdrprimary` command on its successful completion, for example:

```
crtdrsecondary –m QM1 –s 65536 –l myliveappl –i 198.51.100.24 –p 2015
```

Synchronization of data from the main appliance to the recovery appliance begins. (You should preserve the **crtdrsecondary** command in case you need to re-create the secondary queue manager.)

3. Use the **status** command to check the progress of the synchronization, see "status" on page 572.

Configuring queue managers for disaster recovery by using the IBM MQ Console

You set up a disaster recovery configuration by setting up a primary and secondary queue manager.

Before you begin

You must have configured a main appliance and a recovery appliance as described in "Configuring the hardware for disaster recovery" on page 144.

About this task

When you configure disaster recovery, you specify that an existing queue manager on your main appliance is the primary queue manager. When that command completes, it outputs a further command that you run on your recovery appliance to create the secondary queue manager. You then run that command on the recovery appliance to create another instance of that queue manager and specify that it is the secondary queue manager.

You must ensure that there is sufficient free memory for the snapshot of the queue manager data that is required for disaster recovery. For example, a queue manager created with the default 64 GB size requires a further 64 GB of free space to be reserved for the snapshot of the queue manager data.

You can use the chart widget to see the amount of free space available. Configure the widget to display **Disk usage - platform wide**, and select **Appliance data - free space** (see "Monitoring system resource usage by using the IBM MQ Console" on page 178).

Procedure

1. On the appliance that is designated as the main appliance, create the primary instance of the queue manager.
   a. Start the IBM MQ Appliance web UI on the main appliance and click **MQ Console**.

b. In the Queue Manager widget, select the queue manager that you want to designate as a primary instance in the disaster recovery configuration.

c. Ensure that the queue manager is stopped. (Click the stop icon in the queue manager widget toolbar, if necessary.)

d. Select **More** > **Disaster Recovery (DR)** > **Create DR Primary**.

e. Specify the name of the appliance that will host the secondary instance of the queue manager.

f. Specify the IP address of the appliance that will host the secondary instance of the queue manager.

g. Specify the port that will be used for data replication on both appliances. The port number must be between 1025 and 9999 (do not use port 2222, it is reserved by the appliance).

h. Click **Create**.

i. Copy the command that is displayed in the **'Create DR secondary' command** field when the creation of the DR primary has completed. You must run this command on the other appliance in the DR pair before the disaster recovery configuration is complete.

j. Start the queue manager on the main appliance.

2. On the appliance that is designated as the recovery appliance, create the secondary instance of the queue manager.

   a. Start the IBM MQ Appliance web UI on the recovery appliance and click **MQ Console**.

   b. In the Queue Managers widget, select **More** > **Disaster Recovery (DR)** > **Create DR secondary**.

   c. Paste the command that you copied when you configured a primary instance of the queue manager. Pasting the command automatically completes the **Queue manager**, **Volume size**, **Primary appliance name**, **Primary appliance IP**, and **Port** fields with the required values.

   d. Click **Create**.

3. You can view the Disaster Recovery properties of the primary or secondary queue managers to see the progress of the initial synchronization of primary queue manager data with the secondary instance.

Removing a queue manager from a disaster recovery configuration by using the command line

You can remove a queue manager from a disaster recovery configuration by using the command line.

About this task

If you remove a queue manager from a disaster recovery configuration while it is in the primary role, the queue manager is not deleted.

If you remove a queue manager from a disaster recovery configuration while it is in the secondary role, the queue manager is deleted.

Procedure

- To remove a primary queue manager from a disaster recovery configuration:
  1. Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```

2. Ensure that the queue manager is stopped. Enter the following command if necessary:

   ```
   endmqm qmanager
   ```

3. Enter the following command to remove the queue manager from the disaster recovery configuration.

   ```
   dltdrprimary -m qmanager
   ```

- To remove a secondary queue manager from a disaster recovery configuration:
  1. Enter the IBM MQ administration mode by entering the following command:

     ```
     mqcli
     ```

  2. Enter the following command to remove the queue manager from the disaster recovery configuration and delete it.

     ```
     dltdrsecondary -m qmanager
     ```

Removing a queue manager from a disaster recovery configuration by using the IBM MQ Console

You can remove a queue manager from a disaster recovery configuration by using the IBM MQ Console.

About this task

If you remove a queue manager from a disaster recovery configuration while it is in the primary role, the queue manager is not deleted.

If you remove a queue manager from a disaster recovery configuration while it is in the secondary role, the queue manager is deleted.

Procedure

- To remove a primary queue manager from a disaster recovery configuration:
  1. Open the console and find the widget that displays the queue manager.
  2. Ensure that the queue manager is stopped. (Click the stop icon in the queue manager widget toolbar, if necessary.)
  3. Select **More** > **Disaster Recovery (DR)** > **Delete DR Primary**.
  4. Click **Delete** to confirm the removal of the queue manager from the disaster recovery configuration.
- To remove a secondary queue manager from a disaster recovery configuration:
  1. Open the console and find the widget that displays the queue manager.
  2. Select **More** > **Disaster Recovery (DR)** > **Delete DR Secondary**.
  3. Click **Delete** to confirm the deletion of the queue manager.

Viewing the status of a disaster recovery queue manager

You can view the status of a queue manager in a disaster recovery configuration by using the **status** command on the command line, or by using the IBM MQ Console.

About this task

The **status** command returns information about the operational status of a specified queue manager in the disaster recovery configuration. The status can include the following information:

- The disaster recovery role of the queue manager (reported as `Primary` or `Secondary`).
- The current disaster recovery status:

  Normal
  The appliances in the disaster recovery configuration are operating normally.

  Synchronization in progress
  This status can mean that initial replication is completing, or there has been a failure of the disaster recovery replication network and the queue manager has switched into synchronization mode to catch up as quickly as possible.

  Partitioned
  Queue manager data on the appliances is out of step, and cannot be automatically resolved. The **makedrprimary** and **makedrsecondary** commands must be used to resolve the situation. When this status is displayed on one of the appliances in a disaster recovery pair, the other appliance might display the **remote appliance unavailable** status, because the connection was lost before it detected the partitioned status.

  Remote appliance(s) unavailable
  The status means that the connection to the other appliance in the disaster recovery configuration has been lost.

  Inactive
  The queue manager is in the secondary role on both appliances.

  Inconsistent
  This status is shown only when the queue manager is in the secondary role and synchronization is in progress. If you use the **makedrprimary** command on a queue manager that is in this state, the queue manager reverts to the snapshot of its data that was taken before it entered the inconsistent state.

  Reverting to snapshot
  This status is shown when the queue manager is in the secondary role, and the **makedrprimary** command is issued when the queue manager is in the `inconsistent` state. The queue manager is reverted to the current snapshot of its data such that it can run.

  Remote appliance(s) not configured
  This status is shown when the **crtdrprimary** command has been run, to specify that a queue manager has the primary role, but no **crtdrsecondary** command has been run on the other appliance in the disaster recovery pair.
- The percentage complete of a synchronization operation. This information is shown only when the status is `Synchronization in progress`.
- The estimated time at which a synchronization will complete. This information is shown only when the status is `Synchronization in progress`.

- The amount of out-of-sync data that exists on this instance of the queue manager. This is the amount of data written to this instance of the queue manager since it entered the partitioned state. This information is shown only when the status is `Partitioned`.
- The percentage complete of a reversion to snapshot operation. This information is shown only when the status is `Reverting to snapshot`.

Procedure

- To view the DR status of a queue manager by using the command line interface:
  1. Enter the IBM MQ administration mode by entering the following command:
     ```
     mqcli
     ```
  2. View the status of a DR queue manager by entering the following command from one of the appliances:
     ```
     status QMgrName
     ```
     Where:

     QMgrName
     Specifies the name of the DR queue manager that you want to view the status of.
  3. Exit the IBM MQ administration mode by entering the following command:
     ```
     exit
     ```
- To view the DR status of a queue manager by using the IBM MQ Console:
  1. Open the console and find the widget that displays the queue manager.
  2. Select the queue manager in the widget and select the properties icon from the tool bar .
  3. In the properties window, click the **Disaster recovery status** section to open it.

Configuring disaster recovery for a high availability queue manager

You can specify that a high availability queue manager also belongs to a disaster recovery configuration.

Both appliances in a high availability pair are typically located in the same data center. If some disaster befalls the data center, and both appliances are unavailable, you can manually start the queue manager on a recovery appliance located in another data center. See “Disaster recovery for a high availability configuration” on page 6 for an overview.

To configure disaster recovery for a high availability queue manager, complete the following steps:

1. On the appliance where your HA queue manager is running, enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. Stop the HA queue manager:
   ```
   endmqm queue_manager
   ```

3. Specify that the queue manager is the primary instance in a disaster recovery configuration and include a floating IP address that can be used by either of the appliances in the HA pair:

   ```
   crtdrprimary –m queue_manager –r RecoveryName –i RecoveryIP
   –p port_number -f floating_IP
   ```

   Where:

   -m ***QMName***
   Specifies the queue manager that you are preparing for participation in a disaster recovery configuration.

   -r ***RecoveryName***
   Specifies the name of the IBM MQ Appliance that is the recovery appliance.

   -i ***RecoveryIP***
   Specifies the IP address of the recovery appliance.

   -p ***port***
   Specifies the port that the data replication listener on each appliance uses.

   -f ***floatingIP***

   The floating IP address is an IPv4 address that is used to replicate queue manager data from whichever HA appliance the queue manager is currently running on to the queue manager on the recovery appliance. The floating IP address must be in the same subnet group as the static IP address assigned to the replication port (eth20) on both appliances.

   Note that you do not physically configure an Ethernet port with this address. Select a free IP address in the same subnet as the replication ports on the two appliances, and specify it in the **crtdrprimary** command to make it the IP used for replication with the recovery appliance. You must specify a different floating IP address for each of the HA queue managers that you configure disaster recovery for.

   The `crtdrprimary` command configures the queue manager on both appliances in the HA pair, and reserves storage for the data snapshot on both appliances.The **crtdrprimary** command returns a **crtdrsecondary** command when it has completed, for example:

   ```
   Queue manager QM3 is prepared for Disaster Recovery replication.
   Now execute the following command on appliance mydrappl:
   crtdrsecondary –m QM3 –s 65536 –l myliveapp3 –i 198.51.100.10 –p 2015
   ```

4. Copy the **crtdrsecondary** command and run it on the recovery appliance. This creates a secondary version of the queue manager, and queue manager data is replicated from the primary queue manager.

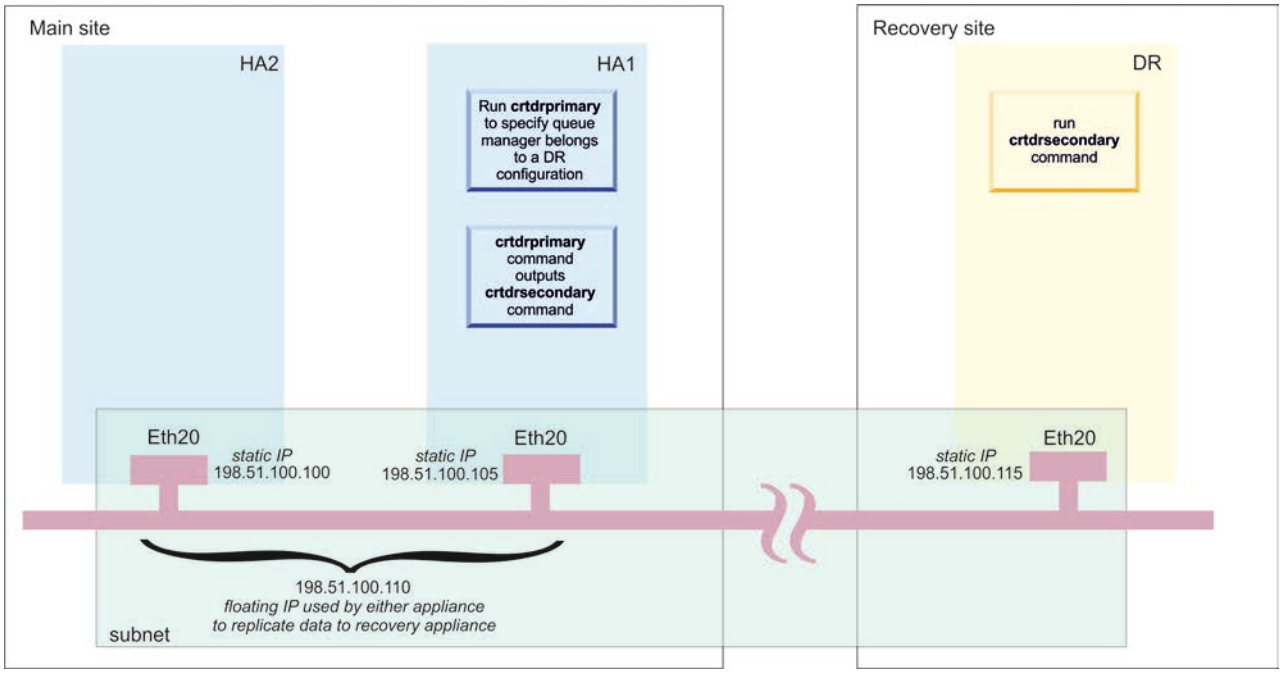

Figure 38. Configuring an HA queue manager for disaster recovery (using example IP addresses)

Example network set up for HA/DR configuration

The example shows the network configuration for an HA pair at the main site with a DR appliance at a recovery site.

The configuration is illustrated in the following diagram. The two HA appliances are located in adjacent racks, and are directly connected to each other with the supplied cables. The replication connection to the recovery appliance and the data connection for applications are made by way of a switch.

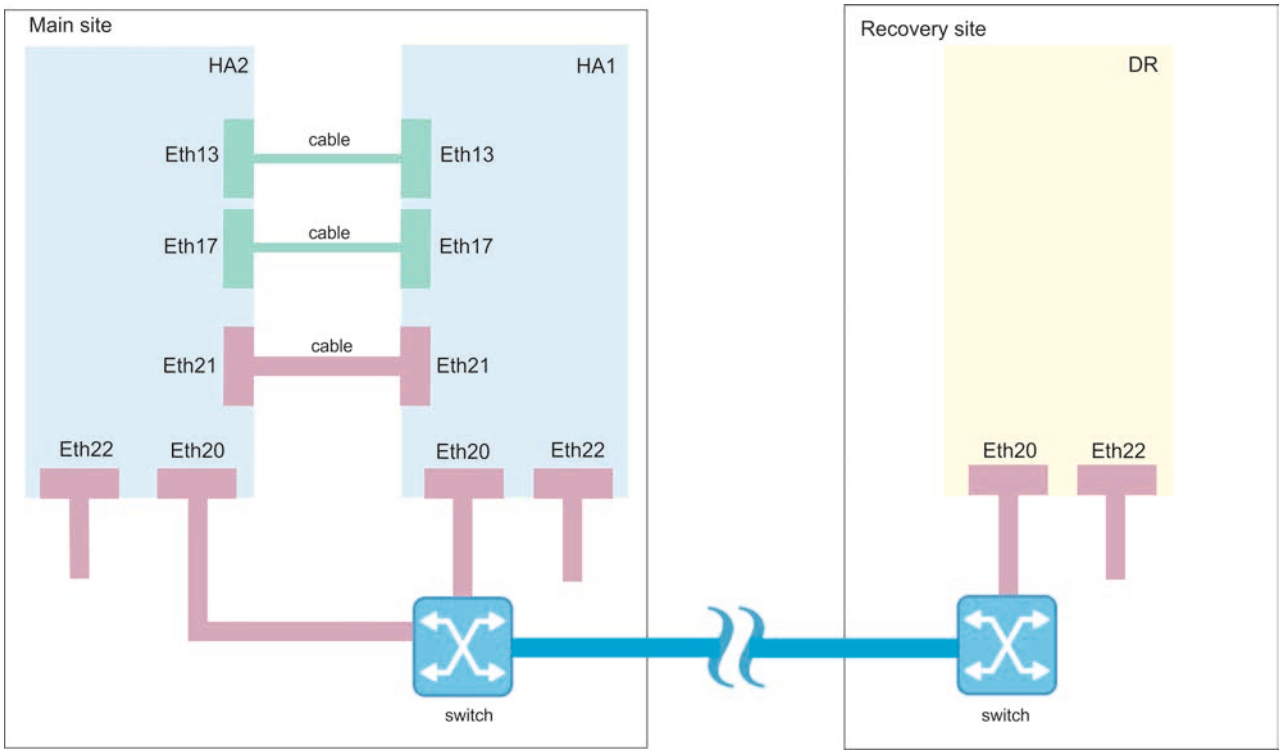

Figure 39. Example HA/DR network configuration

The following tables show how the Ethernet ports are configured on the two HA appliances, HA1 and HA2, and the DR appliance, DR.

Table 18. Ethernet ports on appliance HA1

Ethernet port	IP address	Description
eth13	192.0.10.11	HA primary group interface
eth17	192.0.20.11	HA group alternative interface
eth20	192.0.40.11	Static IP address configured for DR port
-	192.0.40.19	Floating IP address used by appliance HA1 or HA2 for replication with DR appliance
eth21	192.0.30.11	HA replication interface
eth22	203.0.113.0	Data interface

Table 19. Ethernet ports on appliance HA2

Ethernet port	IP address	Description
eth13	192.0.10.12	HA primary group interface
eth17	192.0.20.12	HA group alternative interface
eth20	192.0.40.12	Static IP address configured for DR port

Table 19. Ethernet ports on appliance HA2 (continued)

Ethernet port	IP address	Description
-	192.0.40.19	Floating IP address used by appliance HA1 or HA2 for replication with DR appliance
eth21	192.0.30.12	HA replication interface
eth22	203.0.113.2	Data interface

Table 20. Ethernet ports on appliance DR

Ethernet port	IP address	Description
eth20	192.0.40.13	DR replication interface
eth22	192.0.2.0	Data interface

Chapter 6. Administering

You can administer the IBM MQ Appliance by using the IBM MQ Appliance web UI or by using the command line.

To use the command line to enter IBM MQ commands, you must enter the IBM MQ administration mode. After you enter the IBM MQ administration mode, you can use the control commands and appliance commands that are listed in the Command reference. You can enter the IBM MQ administration mode by using the **mqcli** command. You can exit the IBM MQ administration mode by using the command **exit**.

The following example shows how to enter the IBM MQ administration mode and create a queue manager:

```
mqa# mqcli
mqa(mqcli)# crtmqm QM1
MQ Appliance queue manager created. Creating or replacing default objects for queue manager 'QM1'.
Default objects statistics : 83 created. 0 replaced. 0 failed. Completing setup.
Setup completed.
mqa(mqcli)# exit
mqa#
```

Using the IBM MQ Console

Use the IBM MQ Console to perform common IBM MQ administration tasks.

For the full range of IBM MQ tasks, use the command line interface.

To start the IBM MQ Console:

1. Connect to the IBM MQ Appliance web UI by entering the following URL:

   `https://`*`IP_Address`*`:9090`

   Where *IP_Address* specifies the IP address of the management Ethernet interface (you can determine the IP address of the management Ethernet interface by using the **show int** command).

   Note: This URL uses the default port value. If you changed the port value, replace the `9090` section of the URL with your port number.
2. Click the **MQ Console** icon.

Working with queue managers

You can use the queue manager widget in the IBM MQ Console to create and configure a queue manager. You can then use the widget to control the queue manager.

About this task

The queue manager widget lists the queue managers that have been added to your IBM MQ Appliance. You can select individual queue managers from the list to work with.

You add a queue manager widget to your dashboard by clicking **Add Queue Manager Widget**.

© Copyright IBM Corp. 2015, 2016

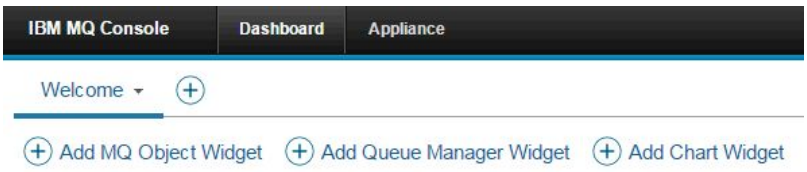

You can configure the queue manager widget by clicking the configure icon in the title bar . You can specify a title for that instance of the widget, and you can specify how many rows of queue managers are displayed per page.

Procedure

- To add a queue manager:
  1. Click the plus icon in the queue manager widget toolbar.
  2. Enter a name for the new queue manager. The name can contain up to 48 characters. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Enter an available TCP/IP port for the queue manager to listen on. The port number must not exceed 65535. If you are configuring a high availability queue manager, the port must be available on both appliances in the high availability group.
  4. Select **High Availability** to specify that the queue manager belongs to a high availability (HA) group. This option is available only if the appliance is part of an HA group.
  5. Optionally select **Specify file system size** and specify the size of the file system created for the queue manager in either GB or MB.
  6. Click **Create**. The new queue manager is created and started.
- To start a queue manager:
  1. Select the queue manager that you want to start from the list in the queue manager widget.
  2. Click the start icon in the queue manager widget toolbar.
- To stop a queue manager:
  1. Select the queue manager that you want to stop from the list in the queue manager widget.
  2. Click the stop icon in the queue manager widget toolbar.
  3. Confirm that you want to stop the queue manager by clicking **Stop**.
- To delete a queue manager:
  1. Select the queue manager that you want to delete from the list in the queue manager widget.
  2. If the queue manager is running, stop it.
  3. Click the delete icon in the queue manager widget toolbar.
  4. Confirm that you want to delete the queue manager by clicking **Delete**. The queue manager and all associated objects are deleted.
- To view and edit the properties of a queue manager:
  1. Ensure that the queue manager is running, and select it in the queue manager list.

2. Click the properties icon in the queue manager widget toolbar.
3. View the properties and click **Edit** to edit them as required. If there is no **Edit** link, the property is read-only, or can be edited only from the command line. Click the help icon to get information about each property, or view the property information in IBM Knowledge Center.

   If the queue manager belongs to a high availability (HA) group, the properties include a **High availability status** category.

- To filter the list of queue managers:
  1. Type your filter text into the filter box.
  2. To stop filtering, click the **X** in the filter box.
- To work with authority records for the queue manager:
  1. Ensure that the queue manager is running, and select it in the queue manager list.
  2. Select one of the following options:
     - Select **More** > **Manage Authority Records** to work with the authority records for the queue manager, and specify what actions groups of users can take.
     - Select **More** > **Manage Create Authority Records** to work with the create authority records for the queue manager, and specify what objects groups of users can create on that queue manager.

     For more information on working with authority records, see working with authority records.
- To view the log for a queue manager:
  1. Select the queue manager. It can be running or stopped.
  2. Select **More** > **View Logs** .
  3. Select the log that you want to view or download, and click the download icon, .
  4. You can view the file, or save it to your local disk, depending on the options offered by your web browser.
- To control the operation of a high availability queue manager:
  1. Ensure that the queue manager is running, and select it in the queue manager list.
  2. Select **More** > **High Availability (HA)** and select one of the following options:
     - **Remove from HA control** - remove the queue manager from the HA group. The queue manager remains as a stand alone queue manager. You can select this option only when the queue manager is running on the current appliance.
     - **Set preferred location** - specify that the current appliance is the preferred location for the queue manager to run on.
     - **Clear preferred location** - specify that the current appliance is no longer the preferred location for the queue manager to run on.
     - **Resolve partitioned data** - after a partitioned state, specify that the queue manager data on the current appliance should be retained, and the data on the other appliance in the HA pair should be discarded.

- To add an existing queue manager to a disaster recovery group:
  1. Ensure that the queue manager is stopped.
  2. Select the queue manager and then select **More** > **Disaster Recovery (DR)** > **Create DR Primary**.
  3. Specify the name of the appliance that will host the secondary instance of the queue manager.
  4. Specify the IP address of the appliance that will host the secondary instance of the queue manager.
  5. Specify the port that will be used for data replication on both appliances. The port number must be between 1025 and 9999 (do not use port 2222, it is reserved by the appliance).
  6. Click **Create**.
  7. Copy the command that is displayed in the **'Create DR secondary' command** field when the creation of the DR primary has completed. You must run this command on the other appliance in the DR pair before the disaster recovery configuration is complete.
- To create a secondary instance of a queue manager in a disaster recovery configuration:
  1. In the Queue Managers widget, select **More** > **Disaster Recovery (DR)** > **Create DR secondary**.
  2. Paste the command that you copied when you configured a primary instance of the queue manager. Pasting the command automatically completes the **Queue manager**, **Volume size**, **Primary appliance name**, **Primary appliance IP**, and **Port** fields with the required values.
  3. Click **Create**.

Working with IBM MQ objects

You use the MQ object widget in the IBM MQ Console to work with the different types of IBM MQ object.

About this task

You add an MQ object widget to your dashboard and then configure it to work with a particular type of IBM MQ object. You then use that instance of the widget to create, configure, and manage objects of that type.

You add an MQ object widget to your dashboard by clicking **Add MQ Object Widget**.

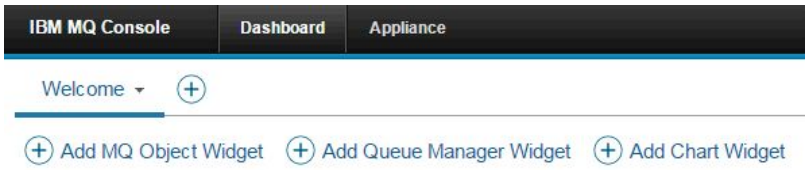

You configure the MQ object widget by clicking the configure icon in the title bar , or by clicking **Choose the configuration options for the object widget**.

To work with other types of IBM MQ objects, that are not available in the console, use MQSC or the IBM MQ Explorer.

Procedure

- To configure the widget to work with queues:
  1. Select the queue manager that the queues will belong to.
  2. Select **Queues** from the **Object** list.
  3. Select the type of queues that you can work with in the widget from the **Object type** list. Choose from:
     - **All**. The widget displays all types of queue.
     - **Local**. The widget displays only local queues. Local queues are physical queues on a local queue manager.
     - **Remote**. The widget displays only remote queues. This queue is not a physical queue - it is the local definition of a queue that exists on a remote queue manager. The local definition of the remote queue contains information that tells the local queue manager how to route messages to the remote queue manager.
     - **Alias**. The widget displays only alias queues. This queue is not a physical queue - it is an alternative name for a local queue. The name of the local queue to which the alias resolves is part of the definition of the alias queue.
     - **Model**. The widget displays only model queues. This queue is not a physical queue - it is a set of queue attributes from which a local queue can be created.

     For more information on queue types, see MQ queues in the IBM MQ documentation.
  4. Select the number of queues to show in each page of the widget from the **Objects per page** list.
  5. Optionally specify a title for this instance of the widget.
  6. Select **Show system objects** to display system queues in the widget.
- To configure the widget to work with topics:
  1. Select the queue manager that the topics will belong to.
  2. Select **Topics** from the **Object** list.
  3. Select the number of topics to show in each page of the widget from the **Objects per page** list.
  4. Specify a title for this instance of the widget.
  5. Select **Show system objects** to display system topics in the widget.
- To configure the widget to work with subscriptions:
  1. Select the queue manager that the subscriptions will belong to.
  2. Select **Subscriptions** from the **Object** list.
  3. Select the number of subscriptions to show in each page of the widget from the **Objects per page** list.
  4. Specify a title for this instance of the widget.
  5. Select **Show system objects** to display system topics in the widget.
- To configure the widget to work with channels:
  1. Select the queue manager that the channels will belong to.
  2. Select **Channels** from the **Object** list.
  3. Select the type of channels that you can work with in the widget from the **Object type** list. Choose from:
     - **All**. The widget displays all types of channel.

- **Cluster-receiver**. The widget displays only cluster-receiver channels. Cluster-receiver channels are used to receive messages in a distributed messaging environment.
- **Cluster-sender**. The widget displays only cluster-sender channels. Cluster-sender channels are used to send messages in a distributed messaging environment.
- **Receiver**. The widget displays only receiver channels. Receiver channels are used to receive messages that are sent to the target queue manager in a point-to-point connection.
- **Requester**. The widget displays only requester channels. Requester channels are used to request a server channel on another queue manager to send messages, or are paired with a sender channel on another queue manager to implement a call-back mechanism.
- **Sender**. The widget displays only sender channels. Sender channels are used to send messages from the source queue manager in a point-to-point connection.
- **Server-connection**. The widget displays only server-connection channels. A server-connection channel is used to connect to an IBM MQ MQI client. The client has a matching client-connection channel to enable bi-directional communication between client and server.
- **Server**. The widget displays only server channels. A server channel is a message channel that the queue manager uses to send messages to other queue managers.

For more information on channel types, see Channels in the IBM MQ documentation.

4. Select the number of channels to show in each page of the widget from the **Objects per page** list.
5. Specify a title for this instance of the widget.
6. Select **Show system objects** to display system channels in the widget.

- To configure the widget to work with client-connection channels:
  1. Select the queue manager that the client-connection channels will belong to.
  2. Select **Client-connection channels** from the **Object** list.
  3. Select the number of client-connection channels to show in each page of the widget from the **Objects per page** list.
  4. Specify a title for this instance of the widget.
  5. Select **Show system objects** to display system objects in the widget. These objects are used internally by IBM MQ and are therefore hidden by default.
- To configure the widget to work with listeners:
  1. Select the queue manager that the listeners will belong to.
  2. Select **Listeners** from the **Object** list.
  3. Select the number of listeners to show in each page of the widget from the **Objects per page** list.
  4. Specify a title for this instance of the widget.
  5. Select **Show system objects** to display system objects in the widget.
- To configure the widget to work with channel authentication records:
  1. Select the queue manager that the channel authentication records will belong to.
  2. Select **Channel Authentication Record** from the **Object** list.

3. Select the number of records to show in each page of the widget from the **Objects per page** list.
4. Specify a title for this instance of the widget.
5. Select **Show system objects** to display system objects in the widget.

- To configure the widget to work with authentication information objects:
  1. Select the queue manager that the authentication information objects will belong to.
  2. Select **Authentication Information** from the **Object** list.
  3. Specify the type of authentication information object to create by selecting one of the following options from the **Object type** list:
     - **CRL LDAP** - Certificate Revocation List checking is done using LDAP servers.
     - **OCSP** - Certificate revocation checking is done using OCSP.
     - **IDPW OS** - Connection authentication user ID and password checking is done using the operating system.
     - **IDPW LDAP** - Connection authentication user ID and password checking is done using an LDAP server.
     - **All** - The type of authentication information object is chosen when the object is created in the widget.
  4. Select the number of records to show in each page of the widget from the **Objects per page** list.
  5. Specify a title for this instance of the widget.
  6. Select **Show system objects** to display system objects in the widget.

Working with queues

After you have configured the MQ object widget, you can use the widget to work with various types of queue.

About this task

You can use the widget to add and delete queues, add and clear messages on a queue, browse messages, and view and set the properties of a queue.

Procedure

- To add a queue:
  1. Click the plus icon in the widget toolbar.
  2. Enter a name for the queue. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Select the type of queue to add.
  4. Click **Create**. The new queue is created.
- To add messages to a queue:
  1. Select the queue that you want to add messages to from the list in the widget.
  2. Click the put message icon
  3. Enter the message that you want to send.
  4. Specify the number of copies of the message that you want to put. Specify a number between 1 and 5000.
  5. Click **Put**.

- To clear messages from a queue:
  1. Select the queue that you want to clear messages from from the list in the widget.
  2. Select **More** > **Clear queue**. Confirm that you want to clear the queue by clicking **Clear Queue**.
- To browse messages in a queue:
  1. Select the queue that you want to browse from the list in the widget.
  2. Click the browse icon . The browse message window opens, displaying messages on the queue.
  3. If you want to have the browser for that queue remain open, click **Add Browse Widget**. You can use the browse widget to view messages, view message properties, add new messages, and clear the queue.
- To delete a queue:
  1. Select the queue that you want to delete from the list in the widget.
  2. If the queue contains messages, clear the messages by selecting **More** > **Clear queue**.
  3. Click the delete icon in the widget toolbar.
  4. Confirm that you want to delete the queue by clicking **Delete**. The queue is deleted.
- To view and edit the properties of a queue:
  1. Select the queue in the widget.
  2. Click the properties icon in the widget toolbar.
  3. View the properties and click **Edit** to edit them. If there is no **Edit** link, the property is read-only or can be edited only from the command line. Click the help icon to get information about each property, or view the property information in IBM Knowledge Center.
- To view and edit authority records for a queue:
  1. Select the queue in the widget.
  2. Click the **More** menu and select **Manage Authority Records**. The authority records show the permissions that users and administrators have on the selected queue. For details of editing the authority records, see working with authority records.

Browsing messages:

You can browse messages by using the browse widget.

About this task

Use the browse widget to view messages, add new messages and clear the queue. You open the browse widget from the queue widget.

Procedure

- To open the browse widget:
  1. Select the queue that you want to browse messages in from the queue widget.

2. Click the browse icon . The browse message window opens, displaying messages on the queue.
3. Click **Add Browse Widget** to add a browse widget to your dashboard. This is like keeping the browse message window open permanently.

- To add new messages:
  1. Click the put message icon .
  2. Type your message in the **Message** field.
  3. Click **Put**.
- To clear the queue of messages:
  1. Select **More** > **Clear queue**.
  2. Click **Clear Queue**.
- To view the properties of a message:
  1. Select the message in the browse widget.
  2. Click the properties icon .
  3. Click **Close** after you have viewed the properties.

Working with topics

After you have configured the MQ object widget, you can use the widget to work with topics.

About this task

You can use the widget to add and delete topics, publish messages to the topic, and view and set the properties of a topic.

Procedure

- To add a topic:
  1. Click the plus icon in the widget toolbar.
  2. Enter a name for the new topic. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Specify the topic string that you publish messages for the topic to. See IBM Knowledge Center for more information.
  4. Click **Create**. The new topic is created.
- To delete a topic:
  1. Select the topic that you want to delete from the list in the widget.
  2. Click the delete icon in the widget toolbar.
  3. Confirm that you want to delete the topic by clicking **Delete**. The topic is deleted.
- To publish a message to a topic:
  1. Select the topic that you want to publish a message to from the list in the widget.
  2. Click the publish icon in the widget toolbar.
  3. Enter the message that you want to publish.
  4. Specify the topic string that you want to publish the message to.

5. Specify the number of copies of the message that you want to publish. Specify a number between 1 and 5000.
6. Click **Publish**.

- To subscribe to a topic:
  1. Select the topic that you want to subscribe to from the list in the widget.
  2. Click the subscribe icon . The Subscribe to Topic window opens.
  3. Specify the topic string that you want to subscribe to, and click **Subscribe**. The window displays messages that are published to the topic string that you are subscribed to.
  4. Click the refresh icon to check for new messages.
  5. If you want to have the subscribe window for that queue remain open, click **Add Subscribe Widget**.
- To view and edit the properties of a topic:
  1. Select the topic in the widget.
  2. Click the properties icon in the widget toolbar.
  3. View the properties and click **Edit** to edit them. If there is no **Edit** link, the property is read-only or can be edited only from the command line. Click the help icon to get information about each property, or view the property information in IBM Knowledge Center.
- To view and edit authority records for a topic:
  1. Select the topic in the widget.
  2. Click the **More** menu and select **Manage Authority Records**. The authority records show the permissions that users and administrators have on the selected topic. For details of editing the authority records, see working with authority records.

Subscribing to topics:

You can view messages from a topic that you have subscribed to by using the subscribe widget.

About this task

You open the subscribe widget from the topic widget.

Procedure

1. Open the topic widget for the topic string that you want to display messages for.
2. Click the subscribe icon .
3. Click **Add Subscribe Widget** to add a subscribe widget to the dashboard.
4. Messages published to the topic string are displayed. You can take the following actions:
   - Refresh the view by clicking the refresh icon:

- Add a name to the widget, and specify how many messages are shown per page by clicking the configure icon:
- Filter the view of messages by entering a string in the **Filter** field.
- View the most recent messages first or last by clicking the arrow in the **Position** column.

Working with subscription objects

After you configure the MQ object widget, you can use the widget to work with subscription objects.

About this task

Subscriptions are issued to a queue manager and contain information about the publications the subscriber wants to receive:

- The topic the subscriber is interested in; this can resolve to multiple topics if wildcards are used.
- An optional selection string to be applied to published messages.
- The name of the queue on which selected publications should be placed.

For more information about subscriptions, see Subscribers and subscriptions and DEFINE SUB in the IBM MQ documentation.

You can also implicitly create a subscription object from a topic widget, and view messages from that subscription. See “Subscribing to topics” on page 166.

Procedure

- To add a subscription object:
  1. Click the plus icon in the widget toolbar.
  2. Specify the name of the object. Valid characters are letters and numbers and the “.”, “/”, “_”, and “%” characters.
  3. Select a **Destination class** of **Managed** or **Provided**. If you select **Managed**, the destination is the local queue manager.
  4. If you select a destination class of **Provided**, in the **Destination** field, specify the name of the queue to which messages for this subscription are forwarded.
  5. In the **Topic string** field, specify the topic string to subscribe to.
  6. Select a **Wildcard usage** setting. Select **Character level wildcard** to specify that wildcard characters represent portions of strings. Select **Topic level wildcard** to specify that wildcard characters represent portions of the topic hierarchy.
  7. Select a **Scope**. Select **All** so the subscription is forwarded to all queue managers directly connected through a publish/subscribe collective or hierarchy. Select **Queue manager** so subscription forwards messages published on the topic only within this queue manager.
  8. Optionally specify a **Selector**. A selection string is an expression that is applied to a publication to determine whether it matches a subscription.
  9. Click **Create**.
- To delete a subscription object:

1. Select the subscription object that you want to delete from the list in the widget.
2. Click the delete icon in the widget toolbar.
3. Confirm that you want to delete the subscription object by clicking **Delete**. The object is deleted.

- To view and edit the properties of a subscription object:
  1. Select the subscription object in the widget.
  2. Click the properties icon in the widget toolbar.
  3. View the properties and click **Edit** to edit them. If there is no **Edit** link, the property is read-only or can be edited only from the command line. Click the help icon to get information about each property.

Working with channels

After you configure the MQ object widget, you can use the widget to work with channels.

About this task

You can use the widget to add and delete channels, start and stop channels, and view and set the properties of a channel.

Procedure

- To add a channel:
  1. Click the plus icon in the widget toolbar.
  2. Enter a name for the channel. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Select the type of channel to add.
  4. If you are creating a sender channel or a server channel, specify the transmission queue that corresponds to the queue manager at the receiver end of the channel.
  5. If you are creating a sender, cluster-sender, or requester channel, specify the connection name. The connection name is the name of the computer that hosts the target queue manager. The format of the name is *computer_name*(*port_number*), where *computer_name* is the name or IP address of the computer that hosts the target queue manager, and *port_number* is the port that the target queue manager's listener is using.
  6. Click **Create**. The new channel is created.
- To delete a channel:
  1. Select the channel that you want to delete from the list in the widget.
  2. Click the delete icon in the widget toolbar.
  3. Confirm that you want to delete the channel by clicking **Delete**. The channel is deleted.
- To start a channel:
  1. Select the channel that you want to start from the list in the queue manager widget.
  2. Click the start icon in the widget toolbar.

- To stop a channel:
  1. Select the channel that you want to stop from the list in the queue manager widget.
  2. Click the stop icon in the widget toolbar.
  3. Confirm that you want to stop the channel by clicking **Stop**.
- To view the properties of a channel:
  1. Select the channel in the widget.
  2. View the properties and click **Edit** to edit them. If there is no **Edit** link, the property is read-only or can be edited only from the command line. Click the help icon to get information about each property, or view the property information in the IBM Knowledge Center.
- To reset a channel:
  1. Select the channel in the widget.
  2. Click the **More** menu and select **Reset**.
  3. Specify a message sequence number. You need to reset a channel if it will not start because the two ends disagree about the sequence number of the next message to send. The message sequence number specifies that number.
  4. Click **Reset Channel**.
- To view or edit authority records for a channel:
  1. Select the channel in the widget.
  2. Click the **More** menu and select **Manage Authority Records**. The authority records show the permissions that users and administrators have on the selected channel. For details of editing the authority records, see working with authority records.

Working with client-connection channels

After you configure the MQ object widget, you can use the widget to work with client-connection channels.

About this task

You can use the widget to add and delete client-connection channels and view and set the properties.

Procedure

- To add a client-connection channel:
  1. Click the plus icon in the widget toolbar.
  2. Enter a name for the client-connection channel. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Specify the connection name. The connection name is the name of the computer that hosts the target queue manager. The format is*computer_name*(*port_number*), where *computer_name* is the name or IP address of the computer that hosts the target queue manager, and *port_number* is the port that the target queue manager's listener is using.
  4. Click **Create**. The new client-connection channel is created.
- To delete a client-connection channel:
  1. Select the client-connection channel that you want to delete from the list in the widget.

2. Click the delete icon in the widget toolbar.
3. Confirm that you want to delete the client-connection channel by clicking **Delete**. The client-connection channel is deleted.

- To view and edit the properties of a client-connection channel:
  1. Select the client-connection channel in the widget.
  2. Click the properties icon in the widget toolbar.
  3. View the properties and click **Edit** to edit them. If there is no **Edit** link, the property is read-only or can be edited only from the command line. Click the help icon to get information about each property, or view the property information in the IBM Knowledge Center.
- To view and edit authority records for a client-connection channel:
  1. Select the client-connection channel in the widget.
  2. Click the **More** menu and select **Manage Authority Records**. The authority records show the permissions that users and administrators have on the selected client-connection channel. For details of editing the authority records, see working with authority records.

Working with listeners

After you configure the MQ object widget, you can use the widget to work with listeners.

About this task

You can use the widget to add and delete listeners, start and stop listeners, view and set listener properties, and view the authority records for a listener.

Procedure

- To add a new listener:
  1. Click the plus icon in the widget toolbar.
  2. Enter a name for the listener. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Click **Create**. The new listener is created.
- To delete a listener:
  1. Select the listener that you want to delete from the list in the widget.
  2. Click the delete icon in the widget toolbar.
  3. Confirm that you want to delete the listener by clicking **Delete**. The listener is deleted.
- To start a listener:
  1. Select the listener that you want to start from the list in the queue manager widget.
  2. Click the start icon in the widget toolbar.
- To stop a listener:
  1. Select the listener that you want to stop from the list in the queue manager widget.

2. Click the stop icon in the widget toolbar.
3. Confirm that you want to stop the listener by clicking **Stop**.

- To view and edit the properties of a listener:
  1. Select the listener in the widget.
  2. Click the properties icon in the widget toolbar.
  3. View the properties and click **Edit** to edit them. If there is no **Edit** link, the property is read-only. Click the help icon to get information about each property, or view the property information in the IBM Knowledge Center.

  Note: By default, listeners are created with the **control** property set to **manual**. You should set it to **Queue Manager** to have it start and stop automatically with the queue manager. On the appliance, listeners always stop automatically when the queue manager stops, whatever the **control** property is set to.
- To view and edit authority records for a listener:
  1. Select the listener in the widget.
  2. Click the **More** menu and select **Manage Authority Records**. The authority records show the permissions that users and administrators have on the selected listener. For details of editing the authority records, see working with authority records.

Working with channel authentication records

After you configure the MQ object widget, you can use the widget to work with channel authentication records.

About this task

To exercise more precise control over the access that is granted to connecting systems at a channel level, you can use channel authentication records.

To enforce security, you can use blocking channel authentication records to block access to your channels. You can also use address map channel authentication records to allow access to specified users. To learn more about channel authentication records, see Channel authentication records in the IBM MQ documentation.

Procedure

- To add a new allowing channel authentication record:
  1. Click the plus icon in the widget toolbar.
  2. Select the **Allow Access** option to indicate that you want to create a rule to allow access to a channel.
  3. Select the identity type that you are allowing from the **Allow this** list. The following types are available:
     - **SSL/TLS subject's Distinguished Name**. Select this option to allow access by a user who is presenting an SSL or TLS personal certificate that contains a specified Distinguished Name.
     - **Client application user ID**. Select this option to allow access by a client application ID from a client-connection channel.

- **Remote queue manager name**. Select this option to allow access by a specified queue manager access.
- **Address**. Select this option to allow specific IP addresses access to all server channels.

4. Specify which channel or channels you are allowing access to by entering a **Channel Profile**. The profile can contain wildcards so that you could block a range of channels. For example, the profile `alphadelta*` blocks channels named alphadelta1, alphadelta2, alphdelta3 and so on.
5. Specify details of what identity is allowed, the field that is displayed depends on the option that is selected for identity type.
   - **Peer Name**. To allow by SSL/TLS subject's Distinguished Name, specify a peer name. For example, `CN=John Smith, O=IBM ,OU=Test , C=GB`. For more details on peer names, see **WebSphere MQ rules for SSLPEER values** in the IBM MQ documentation.
   - **SSL/TLS issuer's Distinguished Name**. To allow by SSL/TLS subject's Distinguished Name, specify the distinguished name of the certificate issuer.
   - **Client user ID**. To allow by client application user ID, specify the user ID of the client that you want to allow.
   - **Remote queue manager**. To allow by remote queue manager name, specify the name of the queue manager whose access you want to allow.
   - **Address**. To allow by address, specify a comma-separated list of IP addresses to be allowed access to all server channels.
6. Click **Create**. The new channel authentication record is created.

- To add a new blocking channel authentication record:

1. Click the plus icon in the widget toolbar.
2. Select the **Block Access** option to indicate that you want to create a rule to block access to a channel.
3. Select the identity type that you are blocking from the **Block this** list. The following types are available:
   - **SSL/TLS subject's Distinguished Name**. Select this option to block any user who is presenting an SSL or TLS personal certificate that contains a specified Distinguished Name.
   - **Client application user ID**. Select this option to block a client application ID from a client-connection channel.
   - **Final assigned user ID**. Select this option to block a list of specified user IDs from a server channel.
   - **Remote queue manager name**. Select this option to block any channel from a specified queue manager.
   - **Address**. Select this option to block specific IP addresses from all server channels.
4. For all types other than **Address**, Specify which channel or channels you are blocking access to by entering a **Channel Profile**. The profile can contain wildcards so that you could block a range of channels. For example, the profile `alphadelta*` blocks channels named alphadelta1, alphadelta2, alphdelta3 and so on.
5. For the **Address** type, specify whether you are matching the address at the listener or the channel level. If you select **Listener**, the address is checked

(and blocked if necessary) before any data is read. If you select **Channel**, the channel that the address is trying to access is discovered before the address is checked.

6. Specify details of what is being blocked, the field that is displayed depends on the option that is selected for identity type.
   - **Peer Name**. To block by SSL/TLS subject's Distinguished Name, specify a peer name. For example, `CN=John Smith, O=IBM ,OU=Test , C=GB`. For more details on peer names, see **WebSphere MQ rules for SSLPEER values** in the IBM MQ documentation.
   - **SSL/TLS issuer's Distinguished Name**. To block by SSL/TLS issuer's Distinguished Name, specify the distinguished name of the certificate issuer.
   - **Client user ID**. To block by client application user ID, specify the user ID of the client that you want to block.
   - **User list**. To block by final assigned user ID, specify a comma-separated list of user IDs to be blocked from server channels.
   - **Remote queue manager**. To block by remote queue manager name, specify the name of the queue manager whose access you want to block.
   - **Address**. To block by address, specify a comma-separated list of IP addresses to be blocked from all server channels.
7. Select **Warning** if you want to run the channel authentication record in warning mode. Warning mode does not block the specified access, but reports access attempts that match the specified rules.
8. Click **Create**. The new channel authentication record is created.

- To delete a channel authentication record:
  1. Select the channel authentication record that you want to delete from the list in the widget.
  2. Click the delete icon in the widget toolbar.
  3. Confirm that you want to delete the channel authentication record by clicking **Delete**. The channel authentication record is deleted.
- To view and edit the properties of a channel authentication record:
  1. Select the channel authentication record in the widget.
  2. Click the properties icon in the widget toolbar.
  3. View the properties and click **Edit** to edit them. If there is no **Edit** link, the property is read-only or can be edited only from the command line. Click the help icon to get information about each property.

Working with authentication information objects

After you configure the MQ object widget, you can use the widget to work with authentication information objects.

About this task

The queue manager authentication information object forms part of IBM MQ support for Transport Layer Security (TLS). These objects contain the definitions

required to perform certificate revocation checking using OCSP or Certificate Revocation Lists (CRLs) on LDAP servers, and the definitions required to enable user ID and password checking.

You can select what type of authentication information object to create when you configure the MQ object widget, or you can select the type **All** and specify the type when you create the actual object.

For more information about authentication information object, see Authentication information objects and DEFINE AUTHINFO in the IBM MQ documentation.

Procedure

- To add a **CRL LDAP** authentication information object:
  1. Click the plus icon in the widget toolbar.
  2. Specify the name of the object. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Specify the LDAP server name. This is the host name, IPv4 dotted decimal address, or IPv6 hexadecimal notation of the host on which the LDAP server is running, with an optional port number.
  4. Click **Create**.
- To add an **OCSP** authentication information object:
  1. Click the plus icon in the widget toolbar.
  2. Specify the name of the object. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Specify the **OCSP responder URL**. This is the URL of the responder used to check for certificate revocation. This value must be an HTTP URL containing the host name and port number of the OCSP responder. If the OCSP responder is using port 80, which is the default for HTTP, then the port number can be omitted. HTTP URLs are defined in RFC 1738.
  4. Click **Create**.
- To add an **IDPW OS** authentication information object:
  1. Click the plus icon in the widget toolbar.
  2. Specify the name of the object. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Click **Create**.
- To add an **IDPW LDAP** authentication information object:
  1. Click the plus icon in the widget toolbar.
  2. Specify the name of the object. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
  3. Specify the LDAP server name. This is the host name, IPv4 dotted decimal address, or IPv6 hexadecimal notation of the host on which the LDAP server is running, with an optional port number. You can specify a comma-separated list of host names.
  4. Specify a field in the LDAP user record to be used as a short user name for this connection.
  5. Click **Create**.
- To delete an authentication information object:

1. Select the authentication information object that you want to delete from the list in the widget.
2. Click the delete icon in the widget toolbar.
3. Confirm that you want to delete the authentication information object by clicking **Delete**. The object is deleted.

- To view and edit the properties of an authentication information object:
  1. Select the authentication information object in the widget.
  2. Click the properties icon in the widget toolbar.
  3. View the properties and click **Edit** to edit them. If there is no **Edit** link, the property is read-only or can be edited only from the command line. Click the help icon to get information about each property.

Working with authority records

You can control the access that groups have to queue managers and IBM MQ objects by specifying an authority record for that group.

About this task

Authority records allow you to fine tune the access that a group of messaging users has to a particular queue manager or MQ object. You configure the authority record in the same way for all object types by using the same procedure, although the actual permissions that you configure depend on the object type.

For example, contrast the different permissions that are available for a queue manager and a queue, as illustrated in the following images:

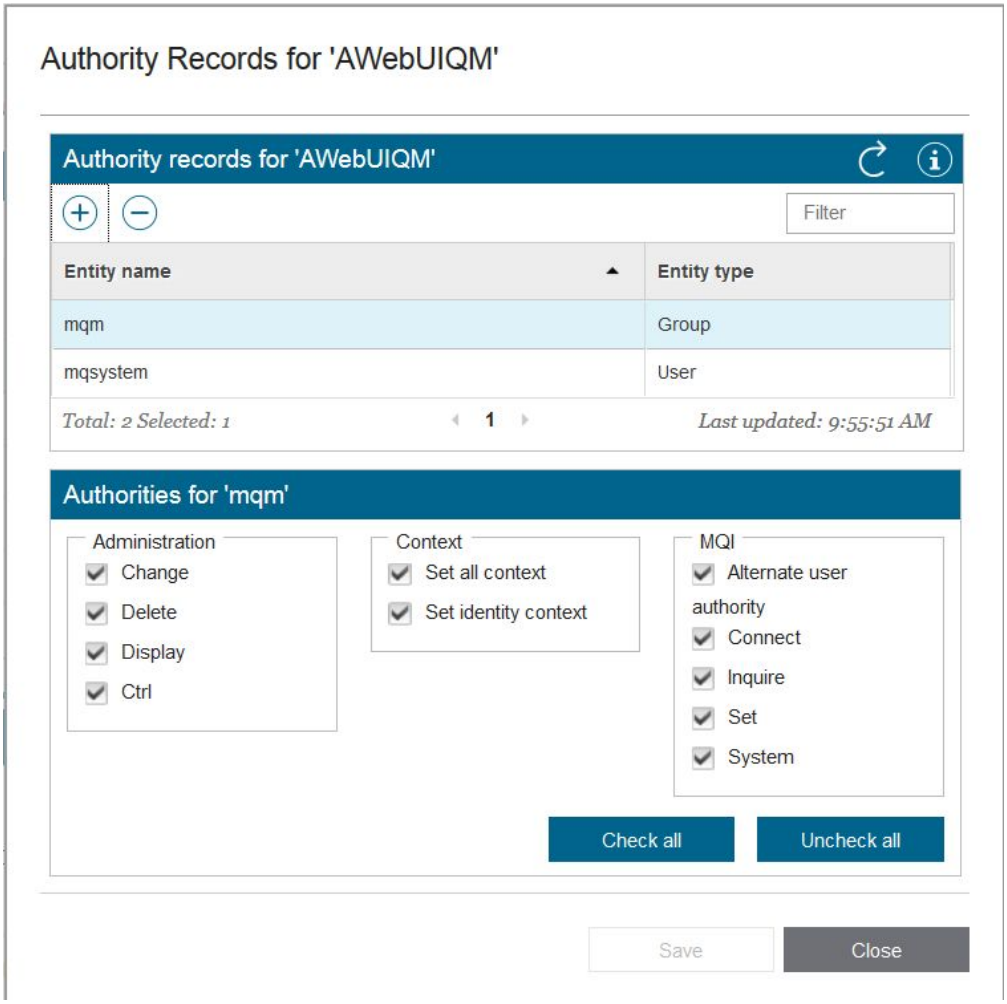
Authority Records for 'AWebUIQM'
Authority records for 'AWebUIQM'
Filter
Entity name
Entity type
mqm
Group
mqsystem
User
Total: 2 Selected: 1
1
Last updated: 9:55:51 AM
Authorities for 'mqm'
Administration
Change
Delete
Display
Ctrl
Context
Set all context
Set identity context
MQI
Alternate user authority
Connect
Inquire
Set
System
Check all
Uncheck all
Save
Close

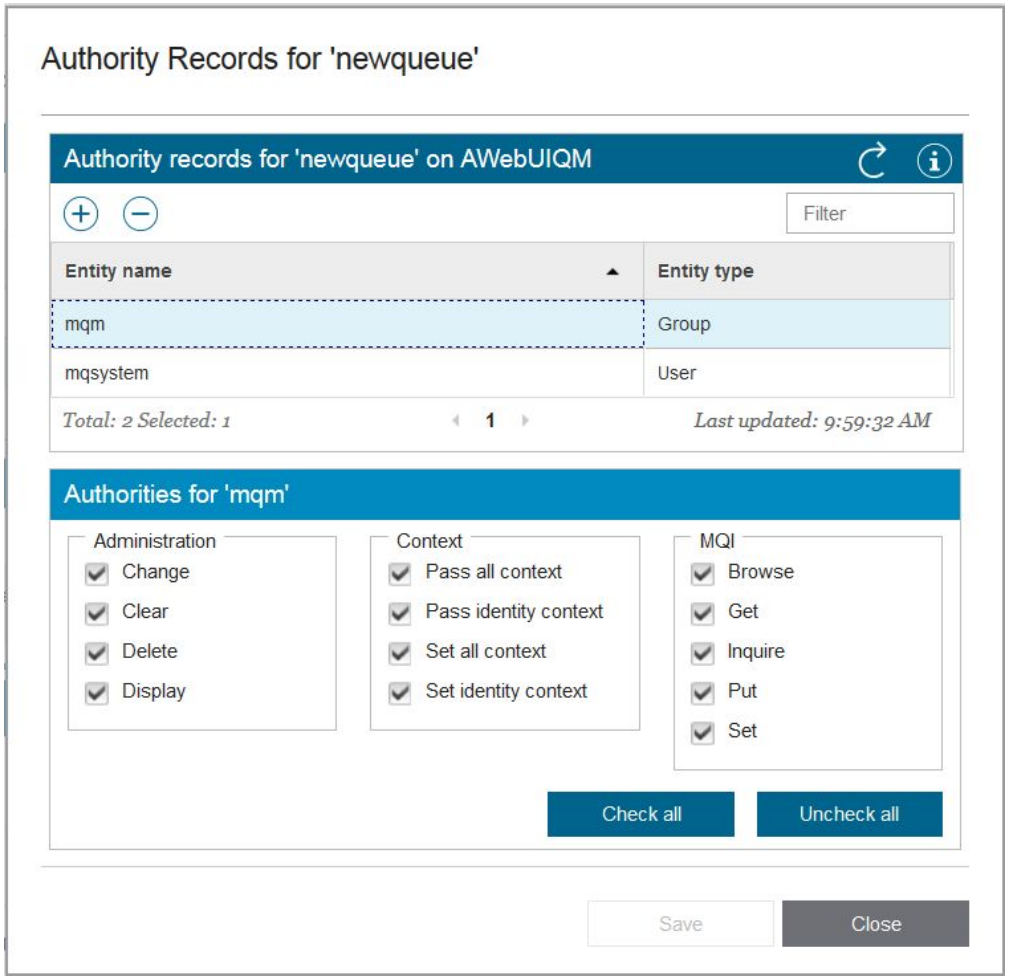

Procedure

- To view or edit an authority record for an MQ object:
  1. Select the object in a widget on the dashboard. The associated queue manager must be running.
  2. From the widget toolbar, select **More** > **Manage Authority Records**.
  3. Select the group that you want to view the authority record for. The authorities for that group are displayed.
  4. Select or deselect authorities as required. Different authorities are available depending on the type of object that you are creating an authority record for.
  5. Click **Save**.
- To view or edit a create authority record for a queue manager:
  1. Select the queue manager in a widget on the dashboard. The queue manager must be running.
  2. From the widget toolbar, select **More** > **Manage Create Authority Records**.
  3. Select the group that you want to view the create authority record for. The authorities for that group are displayed.
  4. Select or deselect create authorities as required.
  5. Click **Save**.
- To create an authority record for an MQ object:

1. Select the MQ object in a widget on the dashboard. The associated queue manager must be running.
2. From the widget toolbar, select **More** > **Manage Authority Records**.
3. Click the plus icon .
4. Select **Group** to specify the type of entity that you want to create an authority record for.
5. Specify the name of the group that you are creating the authority record for. The user or group must already exist.
6. Click **Create**.
7. Select or deselect the authorities that you want the user or group to have. Different authorities are available for each type of object.
8. Click **Save**.

- To create an authority record for creating objects on a queue manager:
  1. Select the queue manager in a widget on the dashboard. The queue manager must be running.
  2. From the widget toolbar, select **More** > **Manage Create Authority Records**.
  3. Click the plus icon .
  4. Select **Group** to specify that you want to create an authority record for a group (IBM MQ Appliance does not support the creation of authorities for individual messaging users).
  5. Specify the name of the group that you are creating the authority record for. The group must already exist.
  6. Click **Create**.
  7. Select or deselect the create authorities that you want the group to have.
  8. Click **Save**.

Monitoring system resource usage by using the IBM MQ Console

You use the Chart widget in the IBM MQ Console to view monitoring data for the IBM MQ Appliance.

About this task

You add a Chart widget to your dashboard and then configure it to monitor a particular aspect of appliance resource usage. You can create many instances of the Chart widget to display different data. The data is displayed in a chart format.

Data is collected at ten-second intervals. The X-axis of the chart displays a timeline. The Y-axis displays units appropriate to the resource that you are viewing. The Y-axis is dynamically resized to accommodate the data that is returned.

You add a Chart widget to your dashboard by clicking **Add Chart Widget**.

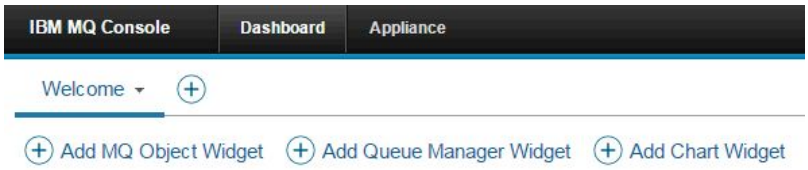

You configure the Chart widget by clicking the configure icon in the title bar . You must have at least one running queue manager before you can configure a chart widget.

After you configure the widget, there is a short delay before data is displayed in the chart. Data is displayed along a time axis as shown in the following example:

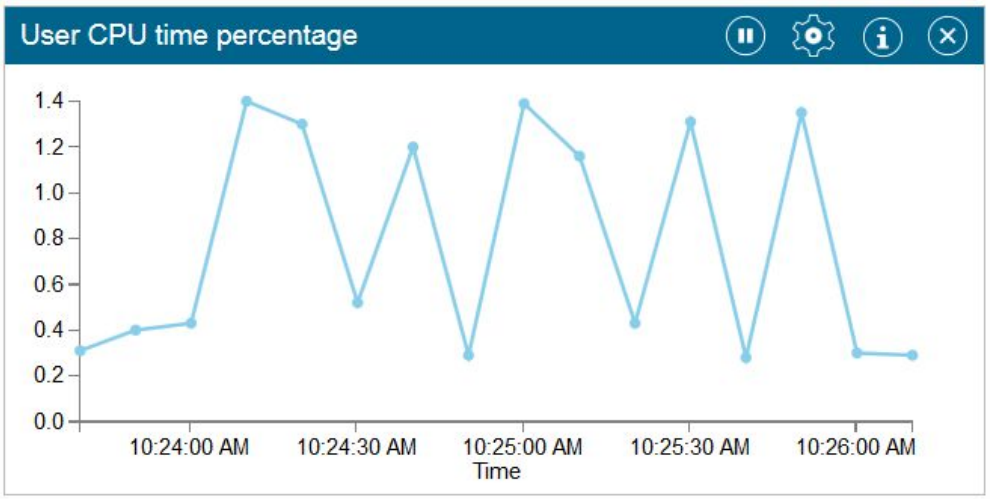

Each data point represents the end of the ten-second period over which the data is collected. You can hover over data points in the chart to see detailed information as shown in the following example:

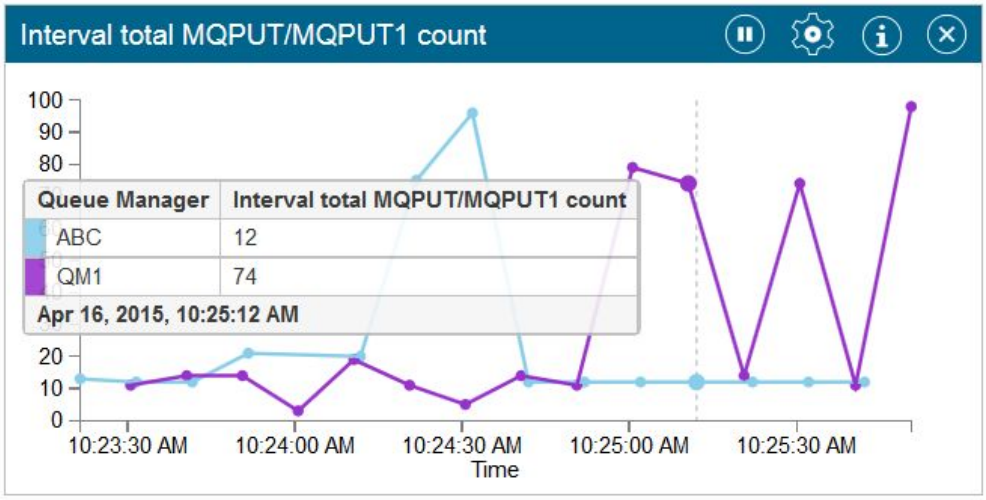

View the video for a demonstration of using the chart widget to monitor a queue manager running on an IBM MQ Appliance.

Using Charts with the IBM MQ Appliance and MQ Console

Procedure

1. Select the class of the resource that you want to monitor:

Platform central processing units
Monitor the usage of the CPUs on the IBM MQ Appliance.

Platform persistent data stores
Monitor the use of disk resource on the IBM MQ Appliance.

API usage statistics
Monitor API calls on the IBM MQ Appliance.

API per-queue usage statistics
Monitor API calls on the IBM MQ Appliance by individual queues. You specify the queue name when you choose this class.

2. Select the type of resource within the chosen class that you want to monitor:

Table 21. Resource types

Class	Type	Description
Platform central processing units	CPU performance – platform wide	Select this type to view performance data for the CPUs and memory in the appliance.
	CPU performance – running queue manager	Select this type to view performance data for the CPUs and memory that is related to the queue managers that you are monitoring. A queue manager must be running for you to monitor it. If you are monitoring results from more than one queue manager, different colors are used to distinguish the performance data in the chart.
Platform persistent data stores	Disk usage – platform wide	Select this type to view performance data for global disk usage within the appliance.
	Disk usage - running queue managers	Select this type to view performance data for the disk usage that is related to the queue managers that you are monitoring. A queue manager must be running for you to monitor it. If you are monitoring results from more than one queue manager, different colors are used to distinguish the performance data in the chart.
	Disk usage - queue manager recovery log	Select this type to view data on how disk storage is currently being used for the recovery log of each queue manager that you are monitoring.
API usage statistics	MQCONN and MQDISC	Select this type to view data on MQCONN and MQDISC calls.
	MQOPEN and MQCLOSE	Select this type to view data on MQOPEN and MQCLOSE calls.
	MQINQ and MQSET	Select this type to view data on MQINQ and MQSET calls.
	MQPUT	Select this type to view data on MQPUT-related calls.
	MQGET	Select this type to view data on MQGET-related calls.

Table 21. Resource types (continued)

Class	Type	Description
	commit and rollback	Select this type to view information about the use of syncpoints by the queue manager.
	subscribe	Select this type to view data related to MQSUB calls.
	publish	Select this type to view data about published messages.
API per-queue usage statistics	MQOPEN and MQCLOSE	Select this type to view data on MQOPEN and MQCLOSE calls for the specified queue.
	MQINQ and MQSET	Select this type to view data on MQINQ and MQSET calls for the specified queue.
	MQPUT and MQPUT1	Select this type to view data on MQPUT-related and MQPUT1-related calls for the specified queue.
	MQGET	Select this type to view data on MQGET-related calls for the specified queue.

3. Select the element for the type of resource monitoring that you selected:

Table 22. Elements for Appliance CPU types

Type	Element	Description
CPU performance – platform wide	User CPU time percentage	Shows the percentage of CPU busy in user state.
	System CPU time percentage	Shows the percentage of CPU busy in system state.
	CPU load – one minute average	Shows the UNIX load average over one minute.
	CPU load – five minute average	Shows the UNIX load average over five minutes.
	CPU load – fifteen minute average	Shows the UNIX load average over fifteen minutes.
	RAM free percentage	Shows the percentage of free RAM memory.
	RAM total bytes	Shows the total bytes of RAM configured.
CPU performance – running queue manager	User CPU time - percentage estimate for queue manager	Estimates the percentage of CPU use in user state for processes related to the queue managers that are being monitored.
	System CPU time - percentage estimate for queue manager	Estimates the percentage of CPU use in system state for processes related to the queue managers that are being monitored.
	RAM total bytes - estimate for queue managers	Estimates the total bytes of RAM in use by the queue managers that are being monitored.

Table 23. Elements for appliance data store resources

Type	Element	Description
Disk usage – platform wide	Appliance data - bytes in use	Shows the overall disk usage on the appliance.
	Appliance data - free space	Shows the overall free space on the appliance.
	System volume - bytes in use	
	System volume - free space	
	MQ trace file system - bytes in use	Shows the number of bytes of disk storage that are being used by the trace file system.
	MQ trace file system - free space	Shows the disk storage that is reserved for the trace file system that is free.
	MQ errors file system - bytes in use	Shows the number of bytes of disk storage that is being used by error data.
	MQ errors file system - free space	Shows the disk storage that is reserved for error data that is currently free.
	MQ FDC file count	Shows the current number of FDC files.
Disk usage - running queue managers	Queue Manager file system - bytes in use	Shows the number of bytes of disk storage that is used by queue manager files for the queue managers that you are monitoring.
	Queue Manager file system - free space	Shows the disk storage reserved for queue manager files that is free.
Disk usage - queue manager recovery log	Log - bytes in use	Shows the number of bytes of disk storage that is used for the recovery logs of the queue managers that you are monitoring.
	Log - bytes max	Shows the maximum bytes of disk storage that are configured to be used for queue manager recovery logs.
	Log file system - bytes in use	Shows the total number of disk bytes in use for the log file system.
	Log file system - bytes max	Shows the number of disk bytes that are configured for the log file system.
	Log - physical bytes written	Shows the number of bytes being written to the recovery logs.
	Log - logical bytes written	Shows the logical number of bytes written to the recovery logs.
	Log - write latency	Shows a measure of the latency when writing synchronously to the queue manager recovery log.

Table 24. Elements for MQI usage resources

Type	Element	Description
MQCONN and MQDISC	MQCONN/MQCONNX count	Shows the number of calls to MQCONN and MQCONNX.
	Failed MQCONN/ MQCONNX count	Shows the number of failed calls to MQCONN and MQCONNX.
	Concurrent connections - high water mark	Shows the maximum number of concurrent connections in the current statistics interval.
	MQDISC count	Shows the number of calls to MQDISC.
MQOPEN and MQCLOSE	MQOPEN count	Shows the number of calls to MQOPEN.
	Failed MQOPEN count	Shows the number of failed calls to MQOPEN.
	MQCLOSE count	Shows the number of calls to MQCLOSE.
	Failed MQCLOSE count	Shows the number of failed calls to MQCLOSE.
MQINQ and MQSET	MQINQ count	Shows the number of calls to MQINQ.
	Failed MQINQ count	Shows the number of failed calls to MQINQ.
	MQSET count	Shows the number of calls to MQSET.
	Failed MQSET count	Shows the number of failed calls to MQSET.
MQPUT	Interval total MQPUT/MQPUT1 count	Shows the number of calls to MQPUT and MQPUT1.
	Interval total MQPUT/MQPUT1 byte count	Shows the total bytes of data that is put by calls to MQPUT and MQPUT1.
	Non-persistent message MQPUT count	Shows the number of non-persistent messages that are put by MQPUT.
	Persistent message MQPUT count	Shows the number of persistent messages that are put by MQPUT.
	Failed MQPUT count	Shows the number of failed calls to MQPUT.
	Non-persistent message MQPUT1 count	Shows the number of non-persistent messages that are put by MQPUT1.
	Persistent message MQPUT1 count	Shows the number of persistent messages that are put by MQPUT1.
	Failed MQPUT1 count	Shows the number of failed calls to MQPUT1.

Table 24. Elements for MQI usage resources (continued)

Type	Element	Description
	Put non-persistent message - byte count	Shows the number of bytes put in non-persistent messages.
	Put persistent message - byte count	Shows the number of bytes put in persistent messages.
	MQSTAT count	Shows the number of calls to MQSTAT.
	Failed MQSTAT count	Shows the number of failed calls to MQSTAT.
MQGET	Interval total destructive get - count	Number of messages removed from queues by MQGET.
	Interval total destructive get - byte count	Bytes of data removed from queues by MQGET.
	Non-persistent message destructive get - count	Number of non-persistent messages removed from queues by MQGET.
	Persistent message destructive get - count	Number of persistent messages removed from queues by MQGET.
	Failed MQGET - count	Shows the number of failed calls to MQGET.
	Got non-persistent messages - byte count	Shows a count of bytes of non-persistent messages returned to MQGET.
	Got persistent messages - byte count	Shows a count of bytes of persistent messages returned to MQGET.
	Non-persistent message browse - count	Shows a count of non-persistent messages that have been browsed.
	Persistent message browse - count	Shows a count of persistent messages that have been browsed.
	Failed browse count	Shows a count of failed message browses.
	Non-persistent message browse - byte count	Shows the number of bytes of non-persistent messages that have been browsed.
	Persistent message browse - byte count	Shows the number of bytes of persistent messages that have been browsed.
	Expired message count	Shows a count of expired messages.
	Purged queue count	Shows a count of queues that have been purged.
	MQCB count	Shows the number of calls to MQCB.

Table 24. Elements for MQI usage resources (continued)

Type	Element	Description
	Failed MQCB count	Shows the number of failed calls to MQCB.
	MQCTL count	Shows the number of calls to MQCTL.
	Failed MQCTL count	Shows the number of failed calls to MQCTL.
Commit and rollback	Commit count	Shows the number of calls to MQCMIT.
	Failed commit count	Shows the number of failed calls to MQCMIT.
	Rollback count	Shows the number of calls to MQBACK.
Subscribe	Create durable subscription count	Shows the number of calls to MQSUB to create durable subscriptions.
	Alter durable subscription count	Shows the number of calls to MQSUB to alter durable subscriptions.
	Resume durable subscription count	Shows the number of calls to MQSUB to resume durable subscriptions.
	Create non-durable subscription count	Shows the number of calls to MQSUB to create non-durable subscriptions.
	Alter non-durable subscription count	Shows the number of calls to MQSUB to alter non-durable subscriptions.
	Resume non-durable subscription count	Shows the number of calls to MQSUB to resume non-durable subscriptions.
	Failed create/alter/resume subscription count	Shows the number of failed calls to MQSUBRQ to create, alter, or resume subscriptions.
	Delete durable subscription count	Shows the number of calls to MQSUB to delete durable subscriptions.
	Delete non-durable subscription count	Shows the number of calls to MQSUB to delete non-durable subscriptions.
	Subscription delete failure count	Shows the number of calls to MQSUB to delete subscriptions.
	MQSUBRQ count	Shows the number of calls to MQSUBRQ
	Failed MQSUBRQ count	Shows the number of failed calls to MQSUBRQ

Table 24. Elements for MQI usage resources (continued)

Type	Element	Description
	Durable subscriber - high water mark	Shows the maximum number of durable subscriptions in the current statistics interval.
	Durable subscriber - low water mark	Shows the minimum number of durable subscriptions in the current statistics interval.
	Non-durable subscriber - high water mark	Shows the maximum number of non-durable subscriptions in the current statistics interval.
	Non-durable subscriber - low water mark	Shows the minimum number of non-durable subscriptions in the current statistics interval.
Publish	Topic MQPUT/MQPUT1 interval total	The number of messages put to topics.
	Interval total topic bytes put	The number of message bytes put to topics.
	Published to subscribers - message count	Shows the number of messages published to subscribers.
	Published to subscribers - byte count	Shows the byte count of messages published to subscribers.
	Non-persistent - topic MQPUT/MQPUT1 count	Shows the number of non-persistent messages put to topics.
	Persistent - topic MQPUT/MQPUT1 count	Shows the number of persistent messages put to topics.
	Failed topic MQPUT/MQPUT1 count	Shows the number of failed attempts to put to a topic.

Table 25. Elements for MQI per-queue usage resources

Type	Element	Description
MQOPEN and MQCLOSE	MQOPEN count	Shows the number of calls to MQOPEN.
	MQCLOSE count	Shows the number of calls to MQCLOSE.
MQINQ and MQSET	MQINQ count	Shows the number of calls to MQINQ.
	MQSET count	Shows the number of calls to MQSET.
MQPUT and MQPUT1	MQPUT/MQPUT1 count	Shows the number of calls to MQPUT and MQPUT1.

Table 25. Elements for MQI per-queue usage resources (continued)

Type	Element	Description
	MQPUT byte count	Shows the total bytes of data that is put by calls to MQPUT and MQPUT1.
	MQPUT non-persistent message count	Shows the number of non-persistent messages that are put by MQPUT.
	MQPUT persistent message count	Shows the number of persistent messages that are put by MQPUT.
	MQPUT1 non-persistent message count	Shows the number of non-persistent messages that are put by MQPUT1.
	MQPUT1 persistent message count	Shows the number of persistent messages that are put by MQPUT1.
	non-persistent byte count	Shows the number of bytes put in non-persistent messages.
	persistent byte count	Shows the number of bytes put in persistent messages.
	queue avoided puts	
	queue avoided bytes	
	lock contention	
MQGET	MQGET count	
	MQGET byte count	
	destructive MQGET non-persistent message count	Number of non-persistent messages removed from the queue by MQGET.
	destructive MQGET persistent message count	Number of persistent messages removed from the queue by MQGET.
	destructive MQGET non-persistent byte count	Shows a count of bytes of non-persistent messages returned to MQGET.
	destructive MQGET persistent byte count	Shows a count of bytes of persistent messages returned to MQGET.
	MQGET browse non-persistent message count	Shows a count of non-persistent messages that have been browsed.
	MQGET browse persistent message count	Shows a count of persistent messages that have been browsed.
	MQGET browse non-persistent byte count	Shows the number of bytes of non-persistent messages that have been browsed.
	MQGET browse persistent byte count	Shows the number of bytes of persistent messages that have been browsed.

Table 25. Elements for MQI per-queue usage resources (continued)

Type	Element	Description
	messages expired	Shows a count of expired messages.
	queue purged count	Shows a count of queues that have been purged.
	average queue time	
	queue time	

4. Select a queue manager to monitor, and specify the color to display information in for that queue manager. Click **Add** to add more queue managers. You can specify up to five queue managers.
5. Optionally, specify a title for your Chart widget.
6. Click **Save**.

Interpreting CPU monitoring information

You can use the CPU monitoring information to see how well utilized your appliance is.

You can view the CPU monitoring information graphically by using the IBM MQ Console, or you can view a text report by using the **status** command. Using the IBM MQ Console gives you an ongoing view.

Two types of CPU statistics are reported: time percentage and CPU load. You can view CPU time percentage information for the appliance as a whole, or for a specified queue manager. CPU load is reported for the appliance as a whole.

Time percentage - appliance wide

The **status** command reports the total CPU percentage used. You can configure the chart widget in the IBM MQ Console to display either percentage of system CPU usage, or percentage of user CPU usage.

Time percentage - specified queue manager

The **status** command reports the total CPU percentage used by the specified running queue manager. You can configure the chart widget in the IBM MQ Console to display either percentage of system CPU usage, or percentage of user CPU usage for the running queue managers. The chart widget plots a line for each running queue manager on the appliance.

CPU load

The system load is a measure of the amount of computational work that a computer system performs. The load average represents the average system load over a period of time. It appears in the form of three numbers that represent the system load during the last one-, five-, and fifteen-minute periods. Each process using or waiting for CPU (the ready queue or run queue) increments the load number by 1. Each process that terminates decrements it by 1.

The IBM MQ Appliance M2000A/M2001A contains 20 cores, and a load average of 20 corresponds to the machine being 100% utilized. A load average of 10 correlates with the appliance being 50% loaded, and a load average of 40 corresponds to the machine having double the work requested that it can process. (The IBM MQ

Appliance M2000B/M2001B contains 6 cores.)

Example `status` output

In the following example, the appliance is 75% CPU utilized, of which the queue manager (PERF0) is using approximately 7%.

```
mqa(mqcli)# status
Memory:                                 16297MB used, 189.1GB total [8%]
CPU:                                    75%
CPU load:                               6.62, 6.19, 5.07
Internal disk:                          786432MB allocated, 2979.5GB total [26%]
System volume:                          5270MB used, 14.7GB allocated [35%]
MQ errors file system:                  175MB used, 2 FDCs, 15.8GB allocated [1%]
MQ trace file system:                   3034MB used, 31.5GB allocated [9%]


mqa(mqcli)# status PERF0
QM(PERF0)                               Status(Running)
CPU:                                    7.29%
Memory:                                 209MB
Queue manager file system:              1230MB used, 63.0GB allocated [2%]
```

Importing and exporting dashboard layouts

You can save a dashboard layout by exporting it. You can then import that layout to the same, or a different, dashboard.

About this task

When you export a dashboard, you create a .json file on your local disk. You can subsequently import the json file to a dashboard to recreate the layout. You can, if required, totally clear an existing dashboard before you import a layout.

Procedure

- To import a dashboard layout:
  1. Open the dashboard.
  2. Click the import icon

. The Import Dashboard Configuration window opens.
  3. Click **Browse** and browse for the location of the file containing your configuration.
  4. Select the **Replace existing dashboard with imported dashboard tabs** option if you want to replace your existing dashboard rather than add to it.
  5. Click **Import**.
- To export a dashboard layout:
  1. Open the dashboard.
  2. Click the export icon

.
  3. Select the **Save file** option and click **OK**.
- To reset the dashboard and clear all the current widgets:
  1. Open the dashboard.
  2. Click the reset dashboard icon

  3. Confirm that you want to delete all your widgets and clear their configurations by clicking **Reset**.

Dashboard

You can use the dashboard to configure your MQ resources.

Select **Add MQ Queue Manager Widget** to create a new widget for queue managers.

Click the **Add** button on the Queue Manager widget to create a new local queue manager.

Select **Add MQ Object Widget** to create a new widget for MQ objects. Open the configuration area for the widget and select the required object type, for example, **Queues**, then click **Save** to populate the widget.

You can create and delete objects by using the toolbar icons. Select the properties icon to view and edit the detailed properties for that object.

Select **Add Chart Widget** to create a new chart widget to monitor queue managers. Open the configuration area for the widget and specify monitoring details, then click **Save** to populate the widget.

You can rearrange the widgets on the dashboard by dragging and dropping them. You can create new dashboard pages by clicking the **New** tab. Changes to the layout are saved automatically.

Message queue control commands

You can control message queues by using MQSC commands that are issued on the IBM MQ Appliance command line.

To use the IBM MQ control commands, you must enter the IBM MQ administration mode by entering the command `mqcli` on the command line. You can exit the IBM MQ administration mode by entering the command `exit`.

The following control commands are available on the IBM MQ Appliance.

Command	Description
“crtmqm” on page 326	Create a queue manager.
“dltmqm” on page 329	Delete a queue manager.
“dmpmqcfg” on page 330	Dump the configuration of a queue manager.
“dspmq” on page 333	Display information about queue managers.
“endmqm” on page 344	Stop a queue manager.
“runmqsc” on page 354	Run MQSC commands on a queue manager.
“strmqm” on page 357	Start a queue manager.

For a full description of commands that are supported on the IBM MQ Appliance, see “Control commands on the IBM MQ Appliance” on page 13.

Administering messaging users

Messaging users work with IBM MQ.

Messaging users perform operations on messaging resources. They can connect to queue managers remotely to send and receive messages. They can be authorized to remotely manage some aspects of queue managers by using client connections such as the IBM MQ Explorer.

Messaging users are distinct from appliance users, who administer the IBM MQ Appliance and configure IBM MQ on the appliance. See "User authentication and authorization" on page 265 for an explanation of the distinction between the two types of user.

After you create messaging users, you must use SET AUTHREC in runmqsc, or use the IBM MQ Console to grant these users access to the required IBM MQ resources.

The appliance reserves the following user IDs for its own use:

- hacluster
- mqm
- mqsystem
- root
- sshd

You cannot create user IDs with these names, or delete, modify, or list these user IDs.

By default, all users belong to the group `users`. You cannot remove users from the `users` group, but you can add them to additional groups.

The appliance reserves the following groups for its own use:

- haclient
- root
- sshd
- utmp

You cannot create groups with these names, or delete or list these groups.

The appliance also provides the standard IBM MQ mqm group. You cannot delete this group, but you can add users to it.

You administer messaging users, and messaging user groups, by using the command line. The commands are run in IBM MQ administration mode, which is entered by typing `mqcli` on the command line. The following table lists the commands that are available:

Command	**Description**
"usercreate" on page 370	Creates user IDs for messaging users on the IBM MQ Appliance.
"userdelete" on page 371	Deletes messaging users.
"usermodify" on page 371	Modifies messaging users
"userlist" on page 372	Lists messaging users, or lists details of a particular user ID.
"groupcreate" on page 372	Adds user groups for messaging users on the IBM MQ Appliance.
"groupdelete" on page 373	Deletes groups for messaging users.

Command	Description
“grouplist” on page 373	Lists groups for messaging users.
“userbackup” on page 373	Backs up messaging users on the IBM MQ Appliance to a file.
“userrestore” on page 374	Restores messaging users on the IBM MQ Appliance from a file to which they were previously backed up.

Using MQSC commands

You can use MQSC interactively on the IBM MQ Appliance.

You cannot create or edit files on the appliance, and you cannot copy existing script files there, and so you cannot run scripts of MQSC commands directly on the appliance.

You can, however, use MQSC interactively to run individual commands. You do so by using the `runmqsc` command from the mqa(mqcli) prompt. See “runmqsc” on page 354 for details.

For details of MQSC commands, see IBM MQ Script (MQSC) commands in the IBM MQ documentation.

You can also run MQSC scripts on appliance message queues by running scripts from remote machines.

See “Differences between queue managers that are running on the IBM MQ Appliance and an IBM MQ installation” on page 18 for specific information about using commands on the appliance.

Using an IBM MQ client

You can use an IBM MQ client to connect to queue managers that are running on the IBM MQ Appliance.

You can download the IBM MQ V8 Clients SupportPac from the following link: MQC8: WebSphere MQ V8 Clients. You can install an IBM MQ client on a Windows or a Linux platform, and then use it to connect to queue managers that are running on the appliance.

For information on installing the IBM MQ client, see Installing a WebSphere MQ client in the IBM MQ documentation.

You can use the client to run commands on queue managers that are running on the appliance.

You can also use the client to run supplied sample programs that support putting and getting messages, and publishing and subscribing to topics. Use these methods if you cannot use the IBM MQ Console.

Setting up a queue manager to accept client connections

Configure your queue manager to securely accept incoming connection requests from an IBM MQ client.

About this task

You must complete this task before you can run commands from a client, or run sample programs to put or get messages, publish or subscribe to topics, or browse message queues.

This procedure requires that you enter MQSC commands. You use the `runmqsc` command on the IBM MQ Appliance to enter MQSC commands interactively. See "runmqsc" on page 354.

Procedure

1. Type `mqcli` to enter IBM MQ administration mode.
2. Obtain a messaging user ID on the appliance your queue manager is running on (see "Administering messaging users" on page 190). This user ID is the authority under which the client connection runs on the queue manager. For example:

   ```
   usercreate -u testuser -p passw0rd
   ```
3. Create and start a queue manager (see "Message queue control commands" on page 190):

   ```
   crtmqm -p port queue_manager_name
   strmqm queue_manager_name
   ```

   For example:

   ```
   crtmqm -p 1440 test1
   strmqm test1
   ```
4. Enter the `runmqsc` command so that you can enter MQSC commands interactively. For example:

   ```
   runmqsc test1
   ```
5. Define a queue to be used by the sample programs. For example:

   ```
   DEFINE QLOCAL(Q)
   ```
6. Define a channel for the sample program to use:

   ```
   DEFINE CHANNEL('channel-name') CHLTYPE(SVRCONN) TRPTYPE(TCP) +
   DESCR('Channel for use by sample programs')
   ```

   For example:

   ```
   DEFINE CHANNEL ('MDB.SVRCONN') CHLTYPE(SVRCONN) TRPTYPE(TCP)
   ```
7. Create a channel authentication rule that allows only the IP address of your client system to use the channel by entering the MQSC command:

   ```
   SET CHLAUTH('channel-name') TYPE(ADDRESSMAP) ADDRESS('client-machine-IP-address') +
   MCAUSER('messaging-user-id')
   ```

   channel-name is the name of your channel.

   client-machine-IP-address is the IP address of your client system.

   messaging-user-id is the user ID you obtained in step 2.

   For example:

   ```
   SET CHLAUTH ('MDB.SVRCONN') TYPE(ADDRESSMAP) ADDRESS(192.0.2.0) MCAUSER('testuser')
   ```
8. Grant access to connect to and inquire the queue manager by entering the following MQSC command:

   ```
   SET AUTHREC OBJTYPE(QMGR) PRINCIPAL('messaging-user-id') +
   AUTHADD(CONNECT, INQ)
   ```

   messaging-user-id is the user ID you obtained in step 2.

   For example:

```
SET AUTHREC OBJTYPE(QMGR) PRINCIPAL('testuser') AUTHADD(CONNECT,INQ)
```

9. For put and get sample programs, grant access to your queue to allow inquiring and the putting and getting messages by the messaging user ID. Enter the following MQSC commands:

   ```
   SET AUTHREC PROFILE('queue-name') OBJTYPE(QUEUE) +
   PRINCIPAL('messaging-user-id') AUTHADD(PUT, GET, INQ, BROWSE)
   ```

   queue-name is the name of your queue.

   messaging-user-id is the user ID you obtained in step 2.

   For example:

   ```
   SET AUTHREC PROFILE('Q') OBJTYPE(QUEUE) PRINCIPAL('testuser') AUTHADD(PUT,GET,INQ,BROWSE)
   ```

10. For publish/subscribe sample programs, grant access to your topic by the messaging user ID to allow publishing and subscribing. Enter the following MQSC commands:

    ```
    SET AUTHREC PROFILE('SYSTEM.BASE.TOPIC') OBJTYPE(TOPIC) +
    PRINCIPAL('messaging-user-id') AUTHADD(PUB, SUB)
    ```

    messaging-user-id is the user ID you obtained in step 2.

    (This command gives *messaging-user-id* access to any topic in the topic tree. Alternatively, you can define a topic object by using **DEFINE TOPIC** and grant accesses only to the part of the topic tree that is referenced by that topic object.)

    For example:

    ```
    SET AUTHREC PROFILE('SYSTEM.BASE.TOPIC') OBJTYPE(TOPIC) PRINCIPAL('testuser') AUTHADD(PUB, SUB)
    ```

11. Set up the following environment variables on your client system:
    - Set the MQSAMP_USER_ID environment variable to identify the user that is running the sample programs, as defined in step 2. On Windows, enter:

      ```
      SET MQSAMP_USER_ID=userID
      ```

      On Linux, enter:

      ```
      export MQSAMP_USER_ID='userID'
      ```

      For example:

      ```
      SET MQSAMP_USER_ID=testuser
      ```

    - Set the MQSERVER environment variable to identify the channel and port that is used for running the sample programs, as defined in step 6. The *ConnectionName* parameter identifies the IP and port of the appliance. On Windows, enter:

      ```
      SET MQSERVER=ChannelName/TransportType/ConnectionName
      ```

      On Linux, enter:

      ```
      export MQSERVER='ChannelName/TransportType/ConnectionName'
      ```

      For example:

      ```
      SET MQSERVER=MDB.SVRCONN/TCP/192.0.2.24(1440)
      ```

What to do next

Your client application can now run the sample programs to put and get messages to a queue, publish and subscribe to a topic, or browse a message queue.

Configuring queue managers and objects by using a client

You can use an IBM MQ V8 client to run commands on an IBM MQ Appliance.

Running commands from a client allows you to configure queue managers and IBM MQ objects by running scripts, rather than by entering individual commands.

See Overview of IBM MQ MQI clients and runmqsc in the IBM MQ documentation.

Putting and getting messages

If you cannot use the IBM MQ Console to put and get messages, you can run sample programs on an IBM MQ client instead.

The sample programs are the accessible method for putting and getting messages for queues that are running on the IBM MQ Appliance. The sample programs are located in *MQ_INSTALLATION_PATH*\Tools\c\samples\bin (Windows) and *MQ_INSTALLATION_PATH*/samp/bin (Linux).

- For information on running sample programs on a client that connects to a queue manager on an IBM MQ Appliance, see "Setting up a queue manager to accept client connections" on page 192.
- For information on using the get sample program, amqsgetc, see The Get sample programs in the IBM MQ documentation.
- For information on using the put sample program, amqsputc, see The Put sample programs in the IBM MQ documentation.

To put messages by using the sample program from the client system command line:

1. Enter the command:

   ```
   amqsputc queue_name queue_manager_name
   ```

   For example:

   ```
   amqsputc Q test1
   ```
2. When prompted, enter the password for the user ID running the sample program (note that the password is displayed in plain text).
3. Type the messages that you want to put on the queue.

To get messages by using the sample program from the client system command line:

1. Enter the command:

   ```
   amqsgetc queue_name queue_manager_name
   ```

   For example:

   ```
   amqsgetc Q test1
   ```
2. When prompted, enter the password for the user ID running the sample program (note that the password is displayed in plain text).
3. Messages from the queue are displayed.

Follow the related link for information on putting and getting messages from the IBM MQ Console.

Publishing and subscribing

If you cannot use the IBM MQ Console to publish and subscribe to topics, you can run sample programs on an IBM MQ client instead.

The sample programs are the accessible method for publishing and subscribing to topics on the IBM MQ Appliance. The sample programs are located in *MQ_INSTALLATION_PATH*\Tools\c\samples\bin (Windows) and *MQ_INSTALLATION_PATH*/samp/bin (Linux). You can run the sample programs that are supplied with the client to publish and subscribe to topics.

- For information on running sample programs on a client that connects to a queue manager on an IBM MQ Appliance, see "Setting up a queue manager to accept client connections" on page 192.
- For detailed information on using the publish/subscribe sample programs, see The Publish/Subscribe sample programs in the IBM MQ documentation.

To publish messages by using the sample program from the client system command line:

1. Enter the command:

   ```
   amqspubc topic_name queue_manager_name
   ```

   For example:

   ```
   amqspubc mytopic test1
   ```

   The publisher connects to the queue manager named test1 and responds with the output, `target topic is mytopic`.
2. Enter the messages that you want to publish to `mytopic`.

To subscribe to the topic by using the sample program:

1. Open another command window and enter the command:

   ```
   amqssubc topic_name queue_manager_name
   ```

   For example:

   ```
   amqssubc mytopic test1
   ```

   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.
2. Start another subscriber in another command window, and watch both subscribers receive publications.

Note: Remember to set the MQSERVER environment variable in each command window, as described in "Setting up a queue manager to accept client connections" on page 192.

Follow the related link for information on publishing and subscribing to topics from the IBM MQ Console.

Browsing a message queue

If you cannot use the IBM MQ Console to browse message queues, you can run sample programs on an IBM MQ client instead.

The sample programs are the accessible method for browsing message queues on the IBM MQ Appliance. The sample programs are located in *MQ_INSTALLATION_PATH*\Tools\c\samples\bin (Windows) and *MQ_INSTALLATION_PATH*/samp/bin (Linux). You can run the sample programs that are supplied with the client to browse message queues.

- For information on running sample programs on a client that connects to a queue manager on an IBM MQ Appliance, see "Setting up a queue manager to accept client connections" on page 192.
- For detailed information on using the browse sample programs, see The Browse sample programs in the IBM MQ documentation.

To browse messages by using the sample program from the client system command line:

1. Enter the command:

   ```
   amqsbcgc queue_name queue_manager_name
   ```

   For example:

   ```
   amqsbcgc Q test1
   ```

2. When prompted, enter the password for the user ID running the sample program (note that the password is displayed in plain text).
3. Information about the message queues is displayed, followed by messages from the queue.

Creating and downloading a CCDT file

You can create a client channel definition table file and download it from the appliance so that a client application can use the table to connect to a queue manager running on the appliance.

Queue managers store client connection channel information in a client channel definition table (CCDT). This information includes authentication rules you have defined for channels on the queue manager. The table is updated whenever a client connection channel is defined or altered.

IBM MQ client applications use the CCDT to determine the channel definitions and authentication information to connect to a queue manager. For a client application to use the CCDT, it must be copied to a file and downloaded from the appliance to the client machine or to a location from where the client can access it.

To copy the CCDT to a file:

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Create a file containing the CCDT for a queue manager by entering the following command:

   ```
   rcrmqobj –m queue_manager –t clchltab
   ```

A file with the name *queue_manager*_AMQCLCHL.TAB is created and can be found in the mqbackup:// URI on the appliance.

You can use the IBM MQ Appliance web UI to download the file, see "Managing files by using the IBM MQ Appliance web UI" on page 233.

Starting and stopping the appliance

You can shut down and restart the appliance by using the command line interface.

Restarting the appliance

You can restart the appliance if you want to clear memory and temporary space. You might restart, for example, before you copy a new firmware image to the appliance.

You shut down and restart the appliance by using the command line interface. Complete the following steps:

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Enter flash mode by typing `config` to enter configuration mode and then typing `flash` on the command line.
3. Type the following command:

   ```
   shutdown reboot
   ```

Shutting down the appliance

You can shut down the appliance without restarting it.

You shut down the appliance by using the command line interface. Complete the following steps:

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Enter flash mode by typing `global` to enter configuration mode and then typing `flash` on the command line.
3. Type the following command:

   ```
   shutdown halt
   ```

Back up and restore

You can back up various features of your IBM MQ Appliance and restore these features to the same or a different appliance, if required.

After you back up an appliance, you must restore it to the same or another appliance that is running the same firmware level.

To back up your IBM MQ Appliance, you back up the following features:

1. Configuration of the IBM MQ Appliance.
2. Messaging users and groups.
3. Key repository.
4. Queue manager configurations.
5. IBM MQ Appliance web UI configurations.
6. Optionally, IBM MQ messages.

You use URIs to copy the backed-up information from the appliance to safe storage. You also use URIs to restore backed-up information to a target appliance. You restore an appliance in the following order.

1. Restore configuration of the IBM MQ Appliance.
2. Restore messaging users and groups.
3. Re-create the queue managers.
4. Restore key repository.
5. Restore queue manager configurations.

6. Restore IBM MQ Appliance web UI configurations.
7. Optionally, restore IBM MQ messages.

Backing up or saving the appliance configuration

You can back up the configuration of the IBM MQ Appliance and restore it to a different appliance if required. You can also save it locally on the appliance.

About this task

You use the **write memory** command to write the current appliance configuration to the `autoconfig.cfg` file. You can then copy the command from a URI on the appliance to secure storage on another system.

Note that the backup that you take is not secure, so sensitive data, such as appliance user IDs and passwords is not included. You must re-create these items manually if you use the backup to restore an appliance.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
2. Log in as a user in the administrators group.
3. Type the following command to enter configuration mode:
   ```
   config
   ```
4. Type the following command to write the current configuration:
   ```
   write memory
   ```
5. When prompted, confirm that you want to overwrite the current `autoconfig.cfg` file.
   ```
   Overwrite existing autoconfig.cfg? y
   ```
6. Optionally, copy the `autoconfig.cfg` file to a location on another system from where you can write it to back up storage:

   To copy the file by using the command line:

   a. Connect to the command line of the appliance as described in "Command line access" on page 97.
   b. Log in to the appliance as an administrator.
   c. Type `config` to enter configuration mode.
   d. Copy the file by typing the following command:
      ```
      copy config:///autoconfig.cfg scp://username@ipaddress/[/]directorypath
      ```

   To copy the file by using the IBM MQ Appliance web UI:

   a. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
   b. Open the config folder.
   c. Click the **autoconfig.cfg** link to save the file to your local system (the exact method for saving the file depends on the type of browser that you use).

Restoring the appliance configuration

You can restore the configuration of an IBM MQ Appliance to the same or to a different appliance.

About this task

If you are restoring to the same appliance, it will have the same IP address and the same name. The first steps are the same as initially configuring the appliance when you first installed it.

If you are restoring to a different appliance, it must be running the same firmware level.

You copy a previously saved `autoconfig.cfg` file to the target IBM MQ Appliance and then restart the appliance.

Procedure

1. Complete the initial configuration as described in "Setting up the initial firmware configuration" on page 54.
2. Copy the `autoconfig.cfg` to the target appliance.

   To copy the file by using the command line interface:

   a. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
   b. Log in as a user in the administrators group.
   c. Type the following command to enter configuration mode:
      ```
      config
      ```
   d. Copy your saved `autoconfig.cfg` to the target appliance:
      ```
      copy scp://username@ipaddress/[/]directorypath config:///autoconfig.cfg
      ```

   To copy the file by using the IBM MQ Appliance web UI:

   a. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
   b. Click **Actions** for the config folder.
   c. Select **Upload files** from the **Actions** menu.
   d. Click **Browse**, and browse for the location of the file on your local system.
   e. Click **Upload** to upload the file to the config directory on the appliance.
3. If you are not already connected to the appliance command line, connect as described in "Command line access" on page 97.
4. Type **`flash`** to enter initialization mode.
5. Shut down and restart the appliance by entering the following command:
   ```
   shutdown reboot
   ```
6. Recreate any appliance user IDs that were previously configured. See "Configuring appliance users" on page 120.

Backing up messaging users

You can back up messaging user accounts and restore them to a different IBM MQ Appliance if required. This back up and restore feature is intended for disaster recovery.

About this task

You use a command to place a copy of the messaging user accounts in a user-accessible file area on the appliance. You then copy that file to a backup store on another system.

Note: You cannot back up appliance users. When you restore a system you must re-create appliance users manually.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
2. Log in as a user in the administrators group.
3. Type **mqcli** to enter IBM MQ configuration mode.
4. Type the following command to back up the messaging users:
   ```
   userbackup -f user_backup_filename
   ```
   Where -f *user_backup_filename* optionally specifies a file name. If you do not specify a file name, the backup is written to a file named `userbackup-`*`date-time`*, for example, `userbackup-20150219-132655`.
5. Type **exit** to leave IBM MQ configuration mode.
6. Copy the user backup file to another system.

   To copy the file by using the command line interface:

   a. Connect to the command line of the appliance as described in "Command line access" on page 97.
   b. Log in to the appliance as an administrator.
   c. Type `config` to enter configuration mode.
   d. Copy the file by typing the following command:
      ```
      copy mqbackup:///user_backup_filename scp://username@ipaddress/[/]directorypath
      ```

   To copy the file by using the IBM MQ Appliance web UI:

   a. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
   b. Open the mqbackup folder.
   c. Click the backup file name link to save the file to your local system (the exact method for saving the file depends on the type of browser that you use).

Restoring messaging users

You can restore messaging user accounts that you previously backed up.

About this task

You copy a file that contains the backed up user accounts to the target appliance. You then use a command to restore the accounts.

Procedure

1. Copy the file containing the backed up user accounts to the appliance:

   To copy the file by using the command line interface:

   a. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
   b. Log in as a user in the administrators group.
   c. Type the following command to enter configuration mode:
      ```
      config
      ```
   d. Copy your saved backup file to the target appliance:
      ```
      copy scp://username@ipaddress/[/]directorypath/filename mqbackup:
      ```
   e. Type **exit** to leave config mode.

To copy the file by using the IBM MQ Appliance web UI:

a. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
b. Click **Actions** for the mqbackup folder.
c. Select **Upload files** from the **Actions** menu.
d. Click **Browse**, and browse for the location of the backup file on your local system.
e. Click **Upload** to upload the file to the mqbackup directory on the appliance.

2. If you are not already connected to the appliance command line, connect as described in "Command line access" on page 97.
3. Type **mqcli** to enter IBM MQ configuration mode.
4. Type the following command to restore the messaging users:

   ```
   userrestore -f user_backup_filename
   ```

Backing up a key repository

You can back up the queue manager key repository and restore it to a different IBM MQ Appliance if required. This back up and restore feature is intended for disaster recovery.

About this task

You use a command to place a copy of the key repository in a file in a user-accessible file area on the appliance. You then copy that file to a backup store on another system.

The file that contains the queue manager key repository might include private keys. The file is encrypted, but you should take appropriate security precautions when handling the file. You need a password to modify or restore the file, and the password is displayed after file is created. Ensure that you make a note of the password and keep it safe.

You should follow this procedure for every queue manager on your system.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
2. Log in as a user in the administrators group.
3. Type **mqcli** to enter IBM MQ configuration mode.
4. Type the following command to back up the key repository for a queue manager:

   ```
   keybackup -m QmanagerName
   ```

   Where *QmanagerName* specifies the queue manager that you want to back up the key repository for.
5. The appliance displays the following warning:

   ```
   This operation will generate a copy of your queue manager key  repository,
   which may include private keys.  Although encrypted, you  should take appropriate security
   precautions in handling this file.  The  password required if you ever need to modify or
   restore this file will  be displayed after the copy has been created.  Do you wish to continue?
   [Y/N]
   ```

   Enter `Y` to continue.

The command creates a compressed archive (`.tar.gz`) of the key repository files. The archive includes the `.kdb` and `.rdb` files, and the `crl` file, if present. It does not include the password stash file. At completion, the name of the archive file and the password that was stored in the password stash file is displayed. The password is needed to restore the key repository.

6. Type **exit** to leave IBM MQ configuration mode.
7. Type **config** to enter configuration mode.
8. Copy the file containing the backed-up repository to another system.

   To copy the file by using the command line interface:

   a. Connect to the command line of the appliance as described in "Command line access" on page 97.
   b. Log in to the appliance as an administrator.
   c. Type `config` to enter configuration mode.
   d. Copy the file by typing the following command:
      ```
      copy mqbackup:///backup_filename scp://username@ipaddress/[/]directorypath
      ```

   To copy the file by using the IBM MQ Appliance web UI:

   a. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
   b. Open the mqbackup folder.
   c. Click the backup file name link to save the file to your local system (the exact method for saving the file depends on the type of browser that you use).

Restoring a key repository

You can restore a queue manager key repository that you previously backed up.

About this task

You copy a file that contains the archive of a previously backed-up key repository to the target appliance. You then use a command to restore it.

Procedure

1. Copy the file containing the backed up key repository to the appliance:

   To copy the file by using the command line interface:

   a. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
   b. Log in as a user in the administrators group.
   c. Type the following command to enter configuration mode:
      ```
      config
      ```
   d. Copy your saved backup file to the target appliance:
      ```
      copy scp://username@ipaddress/[/]directorypath/filename mqbackup:
      ```
   e. Type **exit** to leave config mode.

   To copy the file by using the IBM MQ Appliance web UI:

   a. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
   b. Click **Actions** for the mqbackup folder.
   c. Select **Upload files** from the **Actions** menu.
   d. Click **Browse**, and browse for the location of the backup file on your local system.

e. Click **Upload** to upload the file to the mqbackup directory on the appliance.

2. If you are not already connected to the appliance command line, connect as described in "Command line access" on page 97.
3. Type **mqcli** to enter IBM MQ configuration mode.
4. Type the following command to restore the key repository to the queue manager:

   ```
   keyrestore -m QmanagerName -file filename -password password
   ```

   Where:

   - *QmanagerName* specifies the queue manager that you want to back up the key repository for.
   - *filename* is the file that contains the key repository archive.
   - *password* is the password that was returned when the key repository archive was created.
5. Use the **listcert** and **detailcert** commands to verify that the contents of the key repository are as expected.

Backing up a queue manager

You can back up the configuration of a queue manager and restore it to a different IBM MQ Appliance.

About this task

You connect to the IBM MQ Appliance by using the command line, and save the queue manager configuration to a file. You use the **dumpmqcfg** command to back up the configuration of a message queue.

For more information, see **Backing up and restoring WebSphere MQ queue manager data** in the IBM MQ documentation.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
2. Log in as a user in the administrators group.
3. Type **mqcli** to enter IBM MQ configuration mode.
4. Type the following command to save the configuration of the queue manager:

   ```
   dmpmqcfg -m QM_name -a
   ```

   Where:

   - *QM_name* is the name of the queue manager that you want to back up.
5. Run the following command to view the qm.ini file for the queue manager, and make a note of any non-standard settings. You must re-create these settings when you restore the queue manager.

   ```
   dspmqini -m QM_name
   ```

   See Queue manager configuration files, qm.ini in the IBM MQ documentation for guidance on interpreting the output from this command.
6. Run the following command to view any environment variables that are set up for the queue manager. You must re-create these variables when you restore the queue manager.

   ```
   dspmqvar -m QM_name
   ```

7. Copy the backed up configuration file to a safe location on another system.

   To copy the file by using the command line interface:

   a. Connect to the command line of the appliance as described in "Command line access" on page 97.
   b. Log in to the appliance as an administrator.
   c. Type `config` to enter configuration mode.
   d. Copy the file by typing the following command:

      ```
      copy mqbackup:///QM_name.cfg scp://yourusername@yourip/[/]yourdirectorypath
      ```

      For example:

      ```
      copy mqbackup:///REGA.cfg scp://midtownjj@server00d//safe/store/REGA.cfg
      ```

   To copy the file by using the IBM MQ Appliance web UI:

   a. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
   b. Open the mqbackup folder.
   c. Click the backup file name link to save the file to your local system (the exact method for saving the file depends on the type of browser that you use).

Restoring a queue manager

You can restore a queue manager by using the **runmqsc** command to run the commands in the output from the **dumpmqcfg** command.

About this task

You create a new queue manager on the target IBM MQ Appliance to restore the backed-up configuration to. You then connect to the IBM MQ Appliance by establishing a client connection and running the restore from the client system.

Procedure

1. Create a new queue manager on the IBM MQ Appliance you want to restore the queue manager to. You can use the IBM MQ Console to create the queue manager, see "Using the IBM MQ Console" on page 157, or you can use MQSC commands. If you want to use MQSC commands, you must complete the following steps:
   a. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
   b. Log in as a user in the administrators group.
   c. Type the following command to open the IBM MQ command line interface:

      ```
      mqcli
      ```

   d. Use the **crtmqm** command to create the new queue manager.
2. If you discovered non-standard settings in the queue manager's `.ini` file when you backed up the queue manager, restore those settings by using the **setmqini** command.
3. If you discovered environment variables when you backed up the queue manager, recreate the variables by using the **setmqvar** name.
4. Restore the key repository as described in "Restoring a key repository" on page 203 and then start the queue manager.
5. Enable a client connection to the target queue manager. You must define and start a TCP listener, define an SVRCONN channel, and allow administrator

access to the queue manager by using this channel. You can use the IBM MQ Console to do this, see "Using the IBM MQ Console" on page 157, or you can use MQSC commands.

6. In a command shell on the client system, define the MQSERVER environment variable to identify the channel that connects to the target queue manager. For example (on a Linux client):

   ```
   export MQSERVER='SYSTEM.ADMIN.SVRCONN/TCP/9.20.233.217(1414)'
   ```

7. On the client system, run the following command to replay on the queue manager the commands that were exported from the backed up queue manager:

   ```
   runmqsc -c QM_name  < QM_file
   ```

   Where *QM_file* is the file containing the queue manager configuration that you created when you backed up the queue manager. You can verify the result of restoring the queue manager configuration in the following ways:

   - Checking the responses from **runmqsc**. Ensure that there were no errors or.FDC files by using the command **dspmqerr -f**.
   - Checking the queue manager error logs for unexpected messages by using the command **dspmqerr -m** *QM_name*.
   - Running **dmpmqcfg** again to dump the object definitions for the restored queue manager and then using a file comparison utility to highlight the differences with the original backup file.

8. If the restored queue manager has sender or receiver channels that connect to other queue managers, you must reset the message sequence numbers that are associated with these channels. On each of the remote queue managers, issue the following command for every channel that connects to the restored queue manager:

   ```
   RESET CHANNEL('channel name')
   ```

9. If the restored queue manager is a member of cluster, it joins the cluster automatically when the cluster channels are re-created. Although the restored queue manager has the same name and IP address as the original, it is a new queue manager with a different QMID. In the cluster, it appears as a new queue manager alongside the original queue manager with that name. In order to purge the details of the original queue manager, on a full repository queue manager, issue the following command:

   ```
   reset cluster(cluster_name) action(forceremove) QMID('original QMID')
   ```

Backing up IBM MQ Appliance web UI configuration data

The configuration data for each user who uses the IBM MQ Console on the IBM MQ Appliance web UI can be copied and backed up.

About this task

The IBM MQ Appliance web UI configuration data for each user is automatically copied to files in a user-accessible area on the appliance. You then copy the files to a backup store on another system.

You should follow this procedure for every user on your system:

Procedure

- To copy the file to another system by using the command line interface:
  1. Connect to the command line of the appliance as described in "Command line access" on page 97.

2. Log in to the appliance as an administrator.
3. Type `config` to enter configuration mode.
4. Copy the file by typing the following command:

   ```
   copy mqwebui:///com.ibm.mq.webui.persistence/username.json scp://yourusername@yourip/[/]your
   ```

   For example:

   ```
   copy mqwebui:///com.ibm.mq.webui.persistence/admin.json scp://midtownjj@server00d//safe/stor
   copy mqwebui:///com.ibm.mq.webui.persistence/billg.json scp://midtownjj@server00d//safe/stor
   ```

- To copy the file by using the IBM MQ Appliance web UI:
  1. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
  2. Open the mqwebui folder.
  3. Click the backup file name link to save the file to your local system (the exact method for saving the file depends on the type of browser that you use).

Restoring IBM MQ Appliance web UI configuration data

You can restore configuration data for each user who uses the IBM MQ Appliance web UI on the IBM MQ Appliance.

About this task

The IBM MQ Appliance web UI configuration data for each user is automatically copied to files in a user-accessible area on the appliance. You can restore these files to the appliance. Complete this task after you create appliance users, and before anyone logs on to use the web ui.

You should follow this procedure for every user whose configuration you want to restore:

Procedure

- To copy the file to the appliance by using the command line interface:
  1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
  2. Log in as a user in the administrators group.
  3. Type the following command to enter configuration mode:

     ```
     config
     ```

  4. Copy your saved user configuration file to the target appliance:

     ```
     copy scp://yourusername@yourip/[/]yourdirectorypath mqwebui:///com.ibm.mq.webui.persistence
     ```

     For example:

     ```
     copy scp://midtownjj@server00d//safe/store/admin.json mqwebui:///com.ibm.mq.webui.persistenc
     copy scp://midtownjj@server00d//safe/store/billg.json  mqwebui:///com.ibm.mq.webui.persisten
     ```

  5. Type **exit** to leave config mode.
- To copy the file to the appliance by using the IBM MQ Appliance web UI:
  1. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
  2. Click **Actions** for the mqwebui folder.
  3. Select **Upload files** from the **Actions** menu.
  4. Click **Browse**, and browse for the location of the user configuration file on your local system.

5. Click **Upload** to upload the user configuration file to the mqwebui directory on the appliance.

What to do next

Your users can log in to the IBM MQ Appliance web UI and check that their layouts are as expected.

Backing up and restoring message data

You might need to back up and restore message data as part of your disaster recovery plan for the IBM MQ Appliance. This need depends on individual applications and their requirements.

About this task

If you can identify a set of queues that contain messages that should be reloaded as part of a recovery procedure, then you can use the **dmpmqmsg** utility to back up this message data, and later to restore it.

You cannot run **dmpmqmsg** on the IBM MQ Appliance. You must run it on a client and connect to the appliance. For instructions on using **dmpmqmsg**, see dmpmqmsg in the IBM MQ documentation.

Factory reset

A factory reset restores the IBM MQ Appliance to its default state.

Be aware that a factory reset deletes all queue managers and messages that are hosted on the appliance. The reset forcibly ends all queue managers and detaches any applications that are connected to them. After the update, you require direct console access to reinitialize the system.

You reinitialize with a firmware image that you download from IBM Fix Central. See “Installing a new level of firmware” on page 93 for instructions on how to download a firmware image.

To do a factory reset, complete the following steps:

1. Connect to the IBM MQ Appliance as described in “Command line access” on page 97. Log in as an administrative user.
2. Enter flash mode by typing `config` to enter configuration mode and then typing `flash` on the command line.
3. Type the following command:

   `reinitialize` *`firmware_image_file`*

   Where *firmware_image_file* specifies the name of the firmware image that is used to reinitialize the appliance. The file must be in the `image:` directory.
4. After the reinitialization is complete, you must log in as `admin`, using the password `admin`, and follow the initial setup procedure described in Initializing the appliance.

Operating in a high availability environment

If you have configured a high availability system, there are certain maintenance functions that you might need to complete.

Suspending an appliance from an HA group for maintenance

When you want to suspend an appliance from a high availability group, for example, to carry out maintenance on the appliance, you perform a managed failover. This procedure transfers all the workload to the remaining appliance in the group.

To achieve the managed failover, you put the appliance that you want to temporarily remove from the group into standby mode. You then resume the appliance after the maintenance is complete.

Note: While you have one appliance in standby mode, your queue managers can run only on the remaining appliance. You should take care to avoid any outage on the second appliance.

You use this technique when you update the firmware on the appliances in your high availability group, for example to apply a fix pack. In this situation, you suspend the first appliance, update the firmware, and then resume it. You can then suspend the other appliance, upgrade the firmware, and then resume it.

Suspending an appliance from an HA group by using the command line

You can temporarily remove an appliance from a high availability group by using the command line interface.

About this task

When you remove an appliance from an HA group, all queue managers that run on the suspended appliance are failed over to the other appliance. You should not suspend both appliances in the HA group at the same time.

Note: If you suspend an appliance in an HA group while a queue manager is synchronizing, the queue manager ends immediately, rather than finishing the synchronization. You can check whether a queue manager is synchronizing by using the **status** command, see "status" on page 572.

Procedure

- To suspend an appliance in a high availability group, complete the following steps:
  1. Log in to the appliance as a user in the administrators group.
  2. Type the following command to enter IBM MQ administration mode:
     ```
     mqcli
     ```
  3. Enter the following command:
     ```
     sethagrp -s
     ```
  4. Enter the following command to check the appliance status:
     ```
     dsphagrp
     ```

     The appliance is shown as being in the `standby` state.
- To resume the appliance after maintenance is finished, complete the following steps:
  1. Log in to the appliance as a user in the administrators group.
  2. Type the following command to enter IBM MQ administration mode:
     ```
     mqcli
     ```
  3. Enter the following command:

```
sethagrp -r
```

4. Enter the following command to check the appliance status:

   ```
   dsphagrp
   ```

   The appliance is shown as being in the `online` state. When the appliance is resumed in the HA group, all queue managers with a preference set for the resumed appliance switch back onto that appliance.

Suspending an appliance from an HA group by using the IBM MQ Console

You can temporarily remove an appliance from a high availability group by using the IBM MQ Console.

About this task

When you remove an appliance from an HA group, all queue managers that run on the suspended appliance are failed over to the other appliance. You should not suspend both appliances in the HA group at the same time.

Note: If you suspend an appliance in an HA group while a queue manager is synchronizing, the queue manager ends immediately, rather than finishing the synchronization. You can check whether a queue manager is synchronizing by using the **`status`** command, see "status" on page 572.

Procedure

1. Start the IBM MQ Appliance web UI and click **MQ Console**.
2. To suspend the appliance, select **Suspend this appliance** from the **High Availability** menu in the console title bar.
3. Click **Suspend** to confirm your action. The status of the appliance changes to **Standby**.
4. To resume the appliance after maintenance is finished, select **Resume this appliance** from the **High Availability** menu in the console title bar.
5. Click **Resume** to confirm your action. The status of the appliance changes to **Online**.

Replacing a failed node in a high availability group

If an appliance that belongs to a high availability (HA) group fails, you can replace the appliance and then restore the HA group by following this procedure.

Before you begin

When a node in an HA group fails, the queue managers fail over to the remaining appliance in the group. To restore high availability function after you replace or repair the failed appliance, you must first deconstruct the HA group by running the queue managers standalone and deleting the HA group from the remaining appliance. You then create a new HA group, and add the queue managers back to it.

Before you create the new group, you must ensure that both appliances are running the same level of firmware. If your new appliance is running a later version of the firmware, you must either upgrade your existing appliance, or downgrade your new appliance.

Procedure

1. On the appliance that did not fail, stop each queue manager by using the following command:

   ```
   endmqm QMname
   ```

2. If the queue manager is part of a disaster recovery configuration as well as part of an HA group, you must remove it from the disaster recovery configuration. Use the following command:

   ```
   dltdrprimary -m QMname
   ```

3. Enter the following command to remove a queue manager from the HA group and run it as a stand-alone queue manager. The queue manager must be stopped before you run this command.

   ```
   sethagrp -e QMname
   ```

   Where *QMname* is the name of the queue manager. The queue manager is removed from the HA group. You can use the **strmqm** command to restart the queue manager and run it in a stand alone configuration while you replace the failed node, if required.

   Repeat this command for all HA queue managers.

4. Delete the HA group by entering the following command:

   ```
   dlthagrp
   ```

5. On both the existing appliance and the replacement appliance, create a new HA group by using the **prepareha** and **crthagrp** commands, as described in "Creating a high availability group" on page 130.

6. On the appliance that did not fail, enter the following command to add a queue manager back to the HA group. The queue manager must be stopped before you run this command.

   ```
   sethagrp -i QMname
   ```

   Where *QMname* is the name of the existing queue manager. The queue manager is added to the group and is started. Repeat for all the queue managers that were previously part of the HA group.

7. Set the preferred appliance for the queue manager by running the following command:

   ```
   sethapreferred QMname
   ```

   Repeat this command for each queue manager. Run the command on the appliance that did not fail if you want that appliance to be the preferred location. Run the command on the replaced or repaired appliance if you want that appliance to be the preferred location.

8. If you want to restore disaster recovery capability to any of the queue managers, follow the instructions in "Configuring disaster recovery for a high availability queue manager" on page 140.

Managing queue manager locations in a high availability group

You can specify that queue managers have a preferred appliance in the high availability pair.

About this task

You can specify that a queue manager always runs on a particular appliance in the HA pair, if that appliance is available. By default, the preferred appliance for a

queue manager is the appliance that the queue manager was created on. You can use the **sethapreferred** command to specify a preferred appliance in circumstances such as replacing a failed node, or specifying the favored appliance when an existing queue manager is added to an HA group. The **sethapreferred** is used in conjunction with the **clearhapreferred** command. You can also use these commands from the IBM MQ Console.

In normal circumstances, when both appliances in the group are available and neither has been suspended, **sethapreferred** can be used to relocate a queue manager immediately from one appliance to another, by executing the command on the target appliance. This might be used to move a queue manager back to its natural home appliance in a controlled manner following an outage triggered automatic failover.

When you issue the **sethapreferred** command, the queue manager immediately ends on its current host appliance and starts on the appliance where the command is issued. To revert to manual control of preferred location (leaving the queue manager running on its new host), issue the **clearhapreferred** command.

If a failure occurs, and the queue manager fails over to the other appliance in the pair, it resumes running on its preferred appliance as soon as that appliance is available. If you have not specified a preferred appliance, in this situation the queue manager continues to run on the appliance it has failed over to unless you manually intervene.

Procedure

- To set the current appliance as the preferred appliance for a queue manager:
  1. Enter the IBM MQ administration mode by entering the following command:
     ```
     mqcli
     ```
  2. Specify that the current appliance is the preferred appliance for the named queue manager:
     ```
     sethapreferred QMName
     ```
- To specify that the current appliance is no longer the preferred appliance for a queue manager:
  1. Enter the IBM MQ administration mode by entering the following command:
     ```
     mqcli
     ```
  2. Specify that the current appliance is no longer the preferred appliance for the named queue manager:
     ```
     clearhapreferred QMName
     ```
- To set the current appliance as the preferred appliance for a queue manager by using the IBM MQ Console:
  1. Select the queue manager in the queue manager widget.
  2. Select **More** > **High Availability (HA)** > **Set preferred location**.
  3. Confirm that you want to set the current appliance as the preferred location for the queue manager.
- To specify that the current appliance is no longer the preferred appliance for a queue manager by using the IBM MQ Console:
  1. Select the queue manager in the queue manager widget.
  2. Select **More** > **High Availability (HA)** > **Clear preferred location**.
  3. Confirm that you no longer want the current appliance to be the preferred location for the queue manager.

Viewing the status of appliances in a high availability group

You can view the status of appliances in a high availability (HA) group by using the command line or the IBM MQ Console.

Viewing the high availability status by using the command line

You can view the status of appliances in a high availability (HA) group by using the **dsphagrp** command on the command line.

About this task

The **dsphagrp** command returns information about the operational status of each of the appliances in the HA group. The status can be one of the following statuses:

- `Online`. The appliance is available.
- `Offline`. The appliance is unavailable.
- `Standby`. The appliance has been temporarily removed from the HA group.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. View the status of the appliances in the HA group by entering the following command from one of the appliances:

   ```
   dsphagrp
   ```

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Viewing the high availability status by using the IBM MQ Console

You can view the status of appliances in a high availability (HA) group by using the IBM MQ Console.

About this task

The appliance can have one of the following statuses:

- `Online`. The appliance is available.
- `Offline`. The appliance is unavailable.
- `Standby`. The appliance has been temporarily removed from the HA group.

Procedure

1. Start the IBM MQ Appliance web UI and click **MQ Console**.
2. Click the High Availability menu in the console title bar. The menu displays the status of both appliances in the group.

Viewing the status of a high availability queue manager

You can view the status of a queue manager in a high availability (HA) group by using the **status** command on the command line , or by using the IBM MQ Console.

About this task

The **status** command returns information about the operational status of a specified queue manager in the HA group. The status can include the following information:

- The high availability role of the queue manager (reported as `Primary` or `Secondary`).
- The current high availability status:

  Normal
  The appliances in the disaster recovery configuration are operating normally.

  This appliance in standby mode
  This status means that the appliance has been suspended (by using the **sethagrp -s** command).

  Secondary appliance in standby mode
  This status means that the other appliance in the HA pair has been suspended (by using the **sethagrp -s** command).

  Both appliances in standby mode
  This status means that both appliances in the HA pair have been suspended (by using the **sethagrp -s** command).

  Secondary appliance unavailable
  This status means that the connections to the other appliance in the HA pair have been lost.

  Partitioned
  If the two appliances lose the replication interface, the queue managers are immediately flagged as partitioned. The running queue manager might accumulate out-of-sync data. The other queue manager remains in standby with no out-of-sync data. When the connection is remade, the partitioned state is automatically resolved and replication is resumed.

  Synchronization in progress
  This status is displayed when the primary queue manager is replicating data to the secondary queue manager.

  Inactive
  The queue manager is inactive on both appliances in the HA pair.
- The preferred appliance setting for the queue manager, set to `This Appliance` or `Other Appliance`.
- The percentage complete of a synchronization operation. This information is shown only when the status is `Synchronization in progress`.
- The estimated time at which a synchronization will complete. This information is shown only when the status is `Synchronization in progress`.
- The amount of out-of-sync data that exists on this instance of the queue manager. This is the amount of data written to this instance of the queue manager since it entered the partitioned state. This information is shown only when the status is `Partitioned`.

Procedure

- To view the HA status of a queue manager by using the command line interface:
  1. Enter the IBM MQ administration mode by entering the following command:
     ```
     mqcli
     ```
  2. View the status of an HA queue manager by entering the following command from one of the appliances:
     ```
     status QMgrName
     ```
     Where:

     QMgrName
     Specifies the name of the HA queue manager that you want to view the status of.
  3. Exit the IBM MQ administration mode by entering the following command:
     ```
     exit
     ```
- To view the HA status of a queue manager by using the IBM MQ Console:
  1. Open the console and find the widget that displays the queue manager.
  2. Select the queue manager in the widget and select the properties icon from the tool bar .
  3. In the properties window, click on the **High availability status** section to open it.

Disaster recovery for a high availability queue manager

Follow this guidance to run a queue manager on a recovery system if both appliances in your high availability pair are unavailable.

You can configure a high availability (HA) queue manager for disaster recovery (DR), see “Configuring disaster recovery for a high availability queue manager” on page 140.

If you can no longer run a queue manager on either appliance in an HA pair (for example, if there is a power outage in the data center) you can start the disaster recovery version of the queue manager at the recovery site. Log in to your recovery appliance, and follow the procedure described in “Switching over to a recovery appliance” on page 216.

After your disaster is resolved, and presuming that your HA appliance pair have been restored or re-created, you can switch the queue manager from running on the recovery appliance back to running on the HA appliance pair. You might have to follow a slightly different procedure depending on whether your data is partitioned and, if so, which version of the data you want to retain. Follow the procedures described in “Switching back to the main appliance” on page 217.

Operating in a disaster recovery environment

There are a number of situations in which you might want to switch over to the other appliance in a disaster recovery configuration.

Disaster recovery

Following the complete loss of the primary queue manager at the main site, you start the secondary queue manager at the recovery site. Applications reconnect to the queue manager at the recovery site and the secondary queue manager processes application messages. The steps taken to revert to the previous configuration depend on the cause of the failure. For example, complete loss of main appliance versus temporary loss.

For steps to take following a temporary loss of the main site, see "Switching over to a recovery appliance" and "Switching back to the main appliance by using the command line" on page 217. For steps to take following permanent failure, see "Replacing a failed node in a disaster recovery configuration" on page 227.

Disaster recovery test support

You can test the disaster recovery configuration by isolating the recovery site and starting the secondary queue manager so that you can test it and ensure that applications can connect to it. Normal processing continues at the main site. After the test you can discard test data and restore the live replication link between main and recovery appliances.

For steps to follow to test the secondary queue manager in isolation, see "Testing the recovery appliance" on page 230.

Reversal of disaster recovery roles

You might require to periodically reverse the roles of the appliances in your disaster recovery configuration. In this scenario, the main appliance becomes the recovery appliance. The primary queue manager is stopped and designated as the secondary queue manager, while the other queue manager on the recovery appliance is started and becomes the primary queue manager. Applications must reconnect to the new primary queue manager.

For steps to follow to reverse disaster recovery roles, see "Reversing disaster recovery roles" on page 230.

Switching over to a recovery appliance

If a disaster occurs in your main site, you take steps to switch over to your recovery site.

About this task

Following the loss of the primary queue manager at the main site, you start the secondary queue manager at the recovery site. Applications reconnect to the queue manager at the recovery site and the secondary queue manager processes application messages.

You can start the secondary queue manager by using the command line or the IBM MQ Console.

Note: This task assumes that the primary queue manager on the main appliance is not running. If the queue manager is running, and is not part of a high availability (HA) configuration, you must stop it by using the **endmqm** command, or by clicking the stop icon in the queue manager widget in the IBM MQ Console. If the queue manager is part of an HA configuration, you must leave it to the HA subsystem to handle the stopping of the queue manager.

Procedure

- To start the secondary queue manager by using the command line interface:
  1. Log in to the recovery appliance as a user in the administrators group.
  2. Type the following command to enter IBM MQ administration mode:
     ```
     mqcli
     ```
  3. Run the following command:
     ```
     makedrprimary -m QMname
     ```
     Where *QMname* is the name of the queue manager.
     If the state of the queue manager is inconsistent when it starts (that is, replication failed from the main site), the queue manager reverts to the previous saved snapshot of its data.
  4. Run the following command to start the queue manager:
     ```
     strmqm QMname
     ```
- To start the secondary queue manager by using the IBM MQ Console:
  1. Start the IBM MQ Appliance web UI and click **MQ Console**.
  2. In the queue manager widget, select the secondary queue manager that you want to make the primary.
  3. Select **More** > **Disaster Recovery (DR)** > **Make DR primary**.
  4. Click **Make DR primary** to confirm that you want to make this queue manager the primary instance.
  5. Click the start icon to start the queue manager .
- Ensure that your applications reconnect to the queue manager on the recovery appliance. Provided that you have defined your channels with a list of alternative connection names, specifying your primary and secondary queue managers, then your applications will automatically connect to the new primary queue manager.

Switching back to the main appliance

When the disaster has been resolved, and the main appliance restored, you can revert to running the queue manager on your main appliance.

Switching back to the main appliance by using the command line

After you have switched operations to a recovery appliance, you can take steps to revert to your main appliance.

Before you begin

If the queue manager that you are restoring to the main appliance is a high availability queue manager, and the data has been partitioned, there are special steps to take before you follow this procedure. It is possible that you have three different versions of the data: one on the HA primary, one on the HA secondary, and another on the recovery appliance in the DR configuration. If this is the case, you must resolve the partitioning in the HA group, and then follow the procedure outlined here to resolve the partitioning between the HA group and the DR

recovery appliance. See “Resolving a partitioned problem in a high availability configuration” on page 318.

About this task

Following recovery from loss of the primary queue manager at the main site, you stop the queue manager at the recovery site and restart it at the main site. The process for switching back depends on the state of the queue manager data. There are three possible states:

- No partitioning has occurred. The data on both data managers is the same.
- The data is partitioned, and you want to retain the data from the queue manager that was running on the recovery appliance.
- The data is partitioned, and you want to retain the data from the original primary queue manager on the main site.

You should follow the procedures outlined below in full, regardless of the current states of your main and recovery queue managers. (You might find that the partitioning actually resolves part way through the procedure, depending on the initial state of your queue managers, but if you try to shorten the procedures, you might find that partitioning fails to resolve at all).

Procedure

- Where no partitioning has occurred:
  1. End the queue manager on the recovery appliance:
     ```
     endmqm QMName
     ```
     Where *QMName* identifies the queue manager.
  2. Specify that the queue manager on the recovery appliance is the secondary queue manager:
     ```
     makedrsecondary -m QMName
     ```
  3. On the main appliance, specify that the original queue manager is now the primary once more:
     ```
     makedrprimary -m QMName
     ```
  4. Start the queue manager on the main appliance:
     ```
     strmqm QMName
     ```

  This sequence is shown in diagrammatic form in Figure 1.

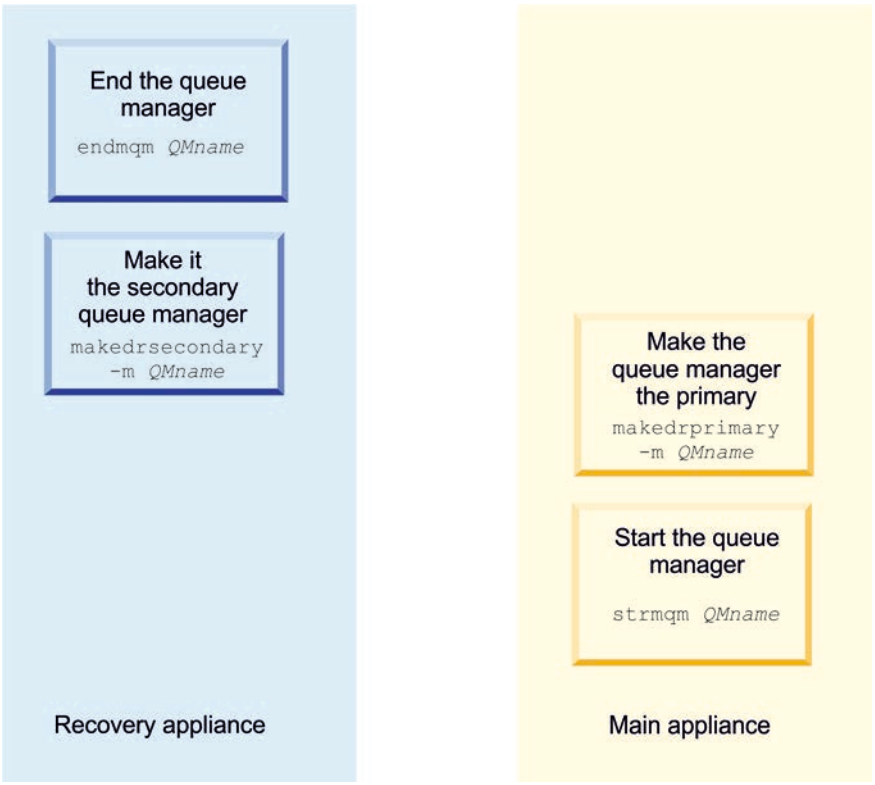

Figure 40. Switching back where there is no partitioning

- Where partitioning has occurred and you want to retain the data from the recovery appliance:
  1. End the queue manager on the recovery appliance:
     ```
     endmqm QMName
     ```
     Where *QMName* identifies the queue manager.
  2. If required, end the queue manager on the main appliance:
     ```
     endmqm QMName
     ```
  3. Specify that the queue manager on the main appliance is the secondary queue manager:
     ```
     makedrsecondary -m QMName
     ```
  4. Specify that the queue manager on the recovery appliance is the primary queue manager:
     ```
     makedrprimary -m QMName
     ```
     Synchronization begins and data is transferred from the recovery appliance to the main appliance.
  5. Check whether the synchronization has completed by using the **status** command:
     ```
     status QMName
     ```
     The output indicates that synchronization is in progress, and displays the percentage complete and estimated completion time.
  6. When the synchronization is complete, specify that the queue manager on the recovery appliance is now the secondary queue manager:
     ```
     makedrsecondary -m QMName
     ```

7. Specify that the queue manager on the main appliance is now the primary queue manager:
   `makedrprimary -m QMName`
8. Start the queue manager on the main appliance:
   `strmqm QMName`

This sequence is shown in diagrammatic form in Figure 2.

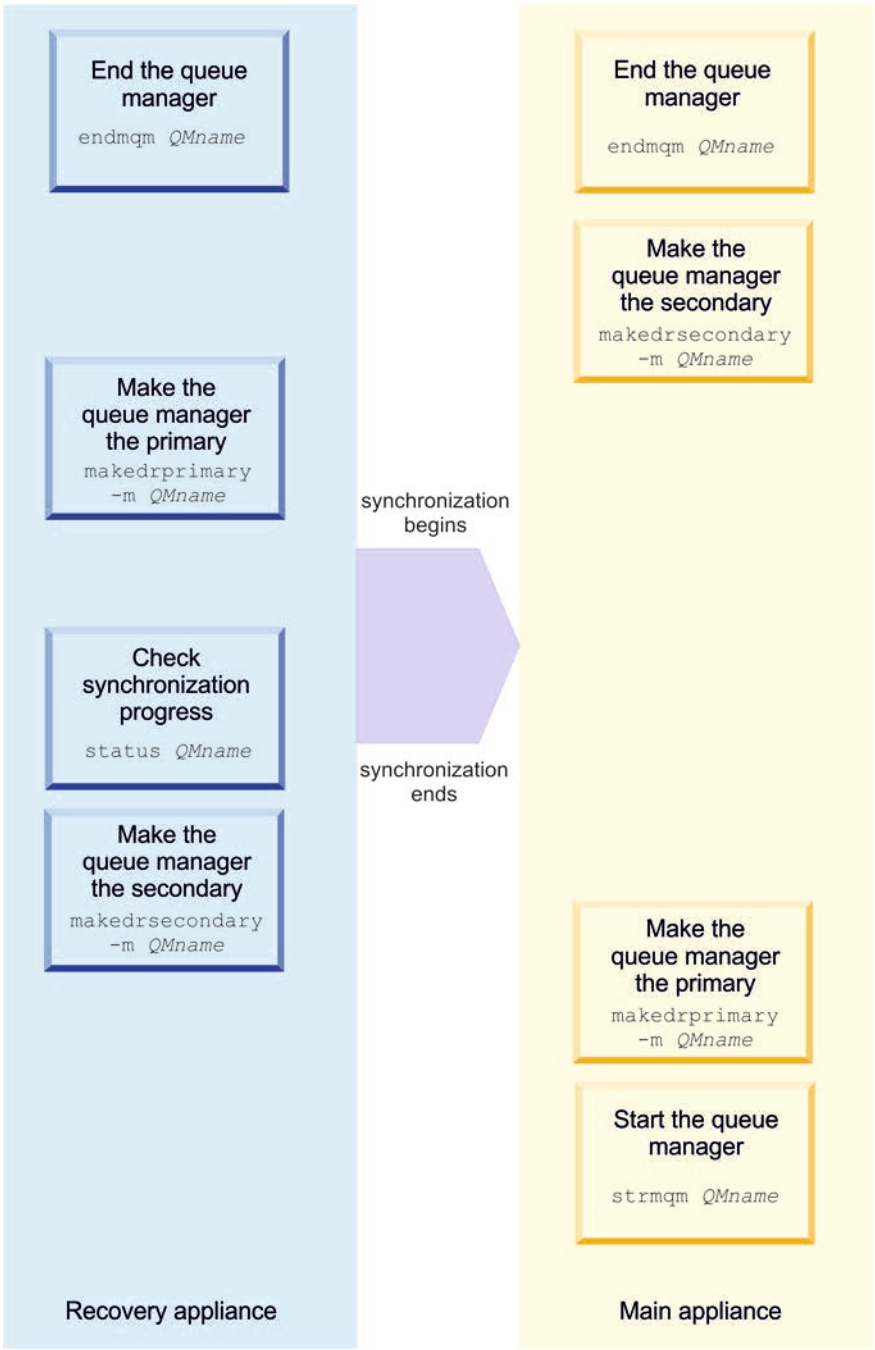

Figure 41. Switching back and retaining recovery appliance data

- Where partitioning has occurred and you want to retain the data from the main appliance:

1. End the queue manager on the recovery appliance:
   `endmqm QMName`

   Where *QMName* identifies the queue manager.
2. If required, end the queue manager on the main appliance:
   `endmqm QMName`
3. Specify that the queue manager on the recovery appliance is now the secondary queue manager:
   `makedrsecondary -m QMName`
4. Specify that the queue manager on the main appliance is now the primary queue manager:
   `makedrprimary -m QMName`

   Synchronization begins and data is transferred from the main appliance to the recovery appliance.
5. Check whether the synchronization has completed by using the **status** command:
   `status QMName`

   The output indicates that synchronization is in progress, and displays the percentage complete and estimated completion time.
6. Start the queue manager on the main appliance (you do not need to wait for synchronization to complete before starting the queue manager):
   `strmqm QMName`

This sequence is shown in diagrammatic form in Figure 3.

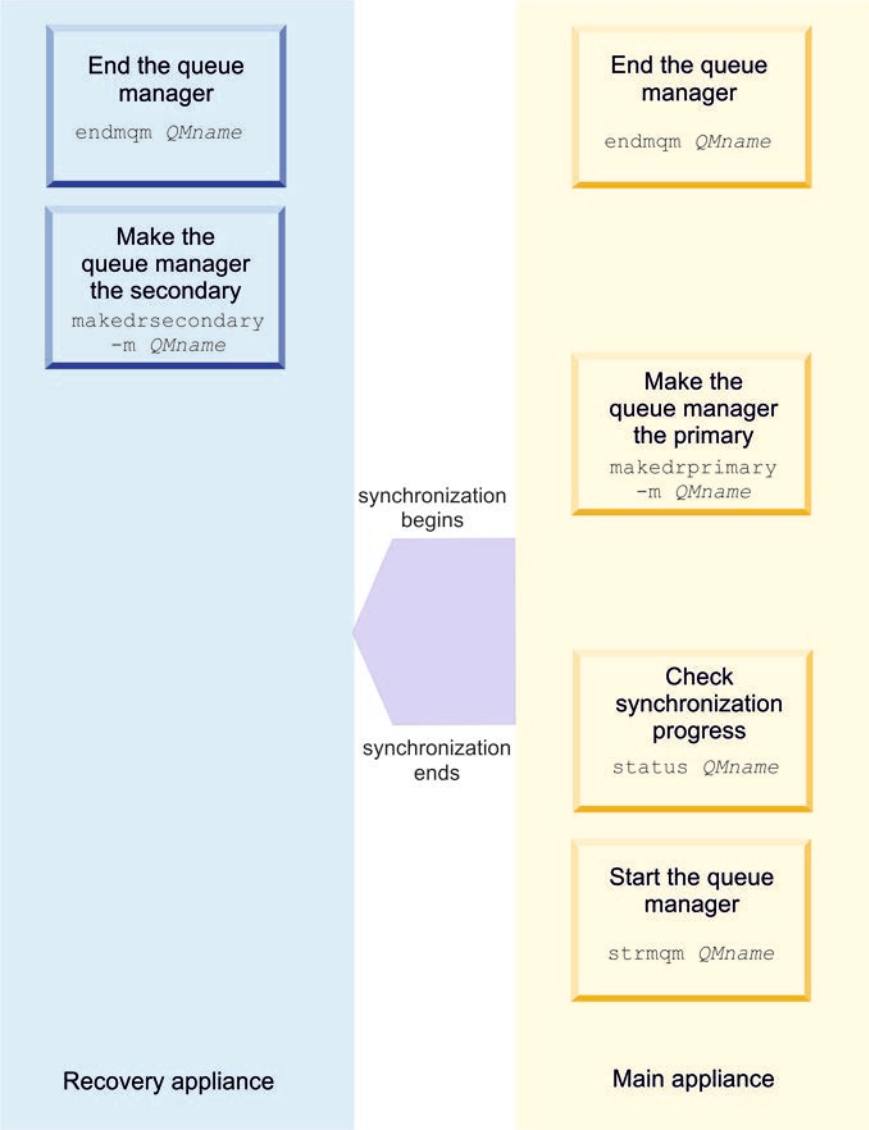

Figure 42. Switching back and retaining main appliance data

Switching back to the main appliance by using IBM MQ Console

After you have switched operations to a recovery appliance, you can take steps to revert to your main appliance.

Before you begin

If the queue manager that you are restoring to the main appliance is a high availability queue manager, and the data has been partitioned, there are special steps to take before you follow this procedure. It is possible that you have three different versions of the data: one on the HA primary, one on the HA secondary, and another on the recovery appliance in the DR configuration. If this is the case, you must resolve the partitioning in the HA group, and then follow the procedure outlined here to resolve the partitioning between the HA group and the DR

recovery appliance. See “Resolving a partitioned problem in a high availability configuration” on page 318.

About this task

Following recovery from loss of the primary queue manager at the main site, you stop the queue manager at the recovery site and restart it at the main site. The process for switching back depends on the state of the queue manager data. There are three possible states:

- No partitioning has occurred. The data on both data managers is the same.
- The data is partitioned, and you want to retain the data from the queue manager that was running on the recovery appliance.
- The data is partitioned, and you want to retain the data from the original primary queue manager on the main site.

You can detect whether partitioning has occurred by viewing the disaster recovery section of the queue manager properties.

You should follow the procedures outlined below in full, regardless of the current states of your main and recovery queue managers. (You might find that the partitioning actually resolves part way through the procedure, depending on the initial state of your queue managers, but if you try to shorten the procedures, you might find that partitioning fails to resolve at all).

Procedure

- Where no partitioning has occurred:
  1. Start the IBM MQ Appliance web UI on the recovery appliance and click **MQ Console**.
  2. In the queue manager widget, end the queue manager on the recovery appliance by selecting it and clicking the stop icon .
  3. Specify that the queue manager on the recovery appliance is the secondary queue manager. Select **More** > **Disaster Recovery (DR)** > **Make DR secondary**.
  4. Start the IBM MQ Appliance web UI on the main appliance and click **MQ Console**.
  5. On the main appliance, specify that the original queue manager is now the primary once more by selecting **More** > **Disaster Recovery (DR)** > **Make DR primary**
  6. Start the queue manager on the main appliance by clicking the start icon

  This sequence is shown in diagrammatic form in Figure 1.

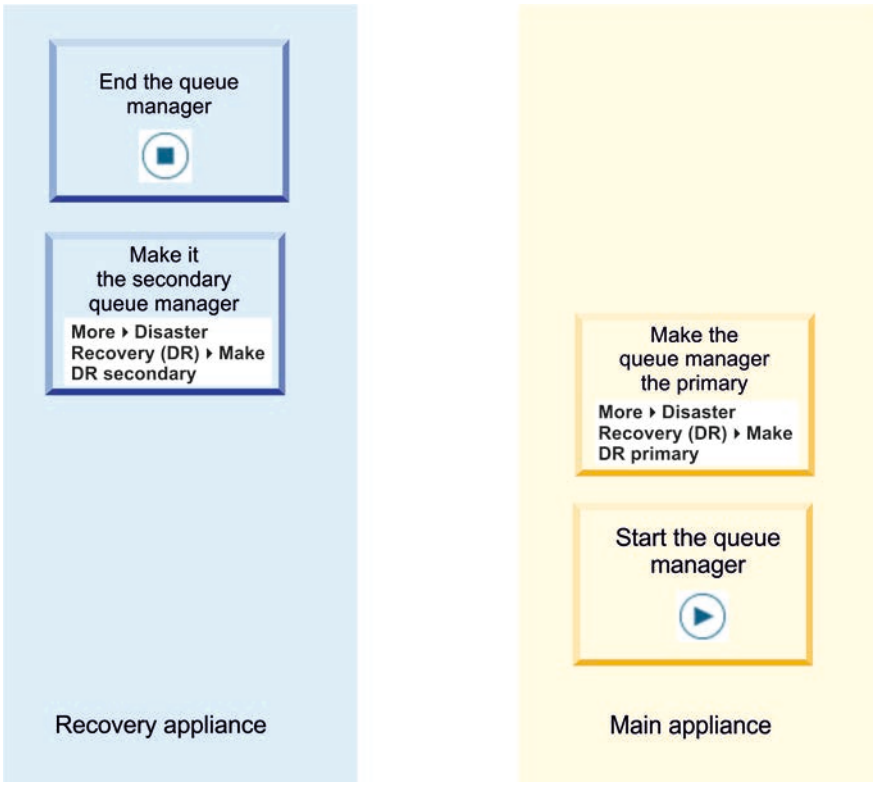

Figure 43. Switching back where there is no partitioning

- Where partitioning has occurred and you want to retain the data from the recovery appliance:
  1. In the queue manager widget, end the queue manager on the recovery appliance by selecting it and clicking the stop icon .
  2. If required, end the queue manager on the main appliance by selecting it and clicking the stop icon .
  3. Specify that the queue manager on the main appliance is the secondary queue manager. Select **More** > **Disaster Recovery (DR)** > **Make DR secondary**.
  4. Specify that the queue manager on the recovery appliance is the primary queue manager. Select **More** > **Disaster Recovery (DR)** > **Make DR primary**. Synchronization begins.
  5. Check whether the synchronization has completed by clicking the properties icon and selecting the **Disaster Recovery** section of the properties.
  6. When the synchronization is complete, specify that the queue manager on the recovery appliance is now the secondary queue manager. Select **More** > **Disaster Recovery (DR)** > **Make DR secondary**.
  7. Specify that the queue manager on the main appliance is now the primary queue manager. Select **More** > **Disaster Recovery (DR)** > **Make DR primary**.
  8. Start the queue manager on the main appliance by clicking the start icon

  This sequence is shown in diagrammatic form in Figure 2.

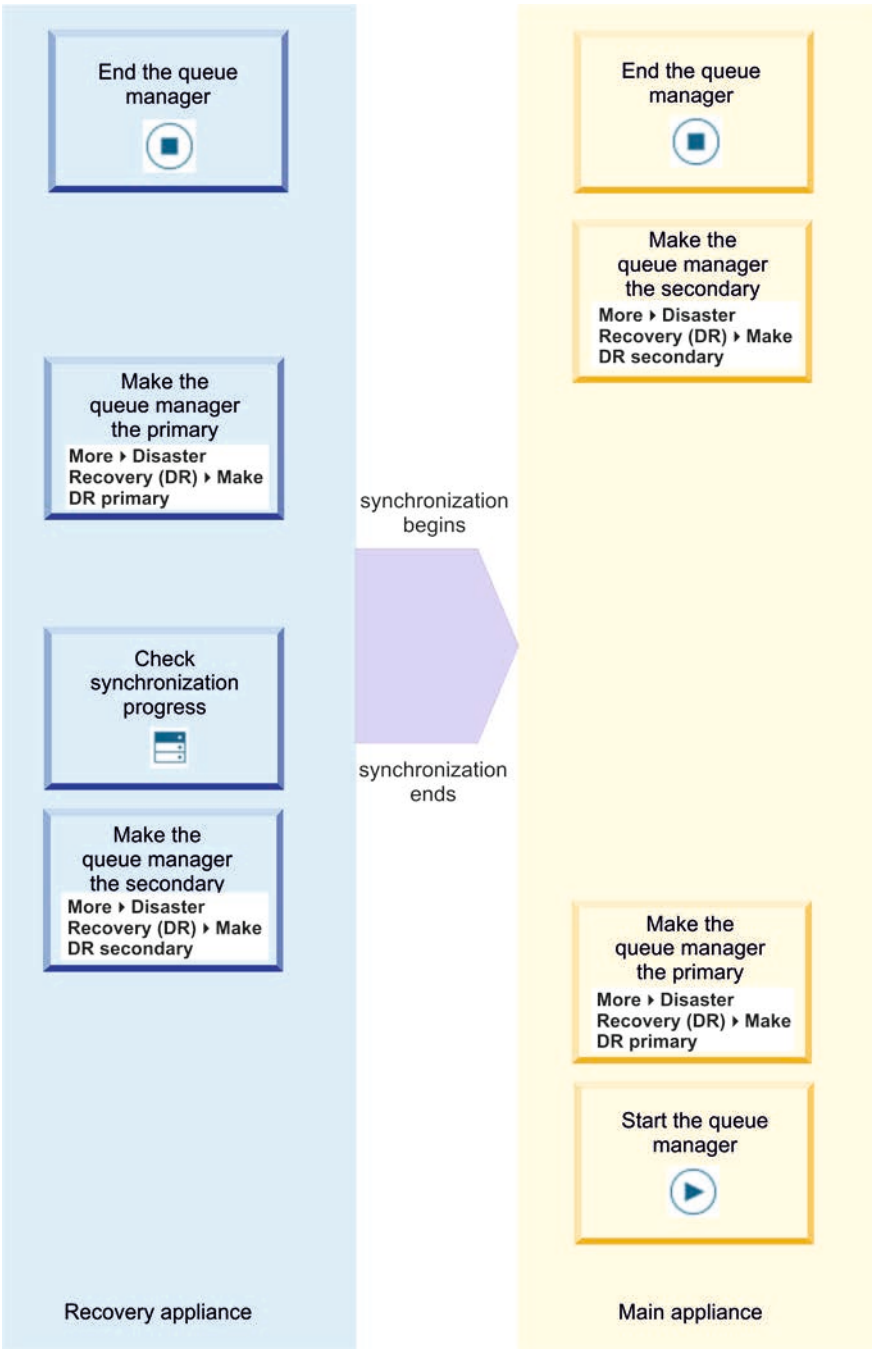

Figure 44. Switching back and retaining recovery appliance data

- Where partitioning has occurred and you want to retain the data from the main appliance:
  1. In the queue manager widget, end the queue manager on the recovery appliance by selecting it and clicking the stop icon .
  2. If required, end the queue manager on the main appliance by selecting it and clicking the stop icon .

3. Specify that the queue manager on the recovery appliance is now the secondary queue manager. Select **More** > **Disaster Recovery (DR)** > **Make DR secondary**.
4. Specify that the queue manager on the main appliance is now the primary queue manager. Select **More** > **Disaster Recovery (DR)** > **Make DR primary**. Synchronization begins.
5. Check whether the synchronization has completed by clicking the properties icon and selecting the **Disaster Recovery** section of the properties.
6. Start the queue manager on the main appliance by clicking the start icon (you do not need to wait for synchronization to complete before starting the queue manager): .

   ```
   strmqm QMName
   ```

This sequence is shown in diagrammatic form in Figure 3.

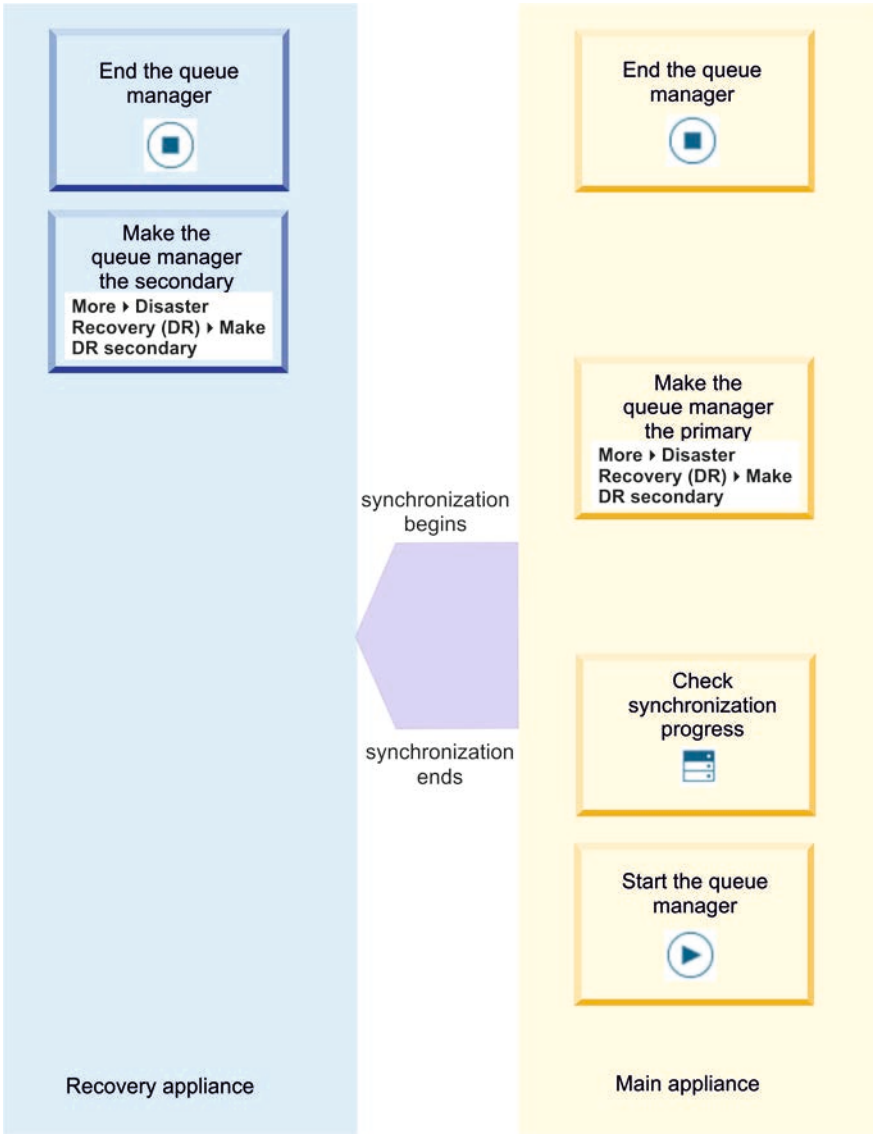

Figure 45. Switching back and retaining main appliance data

Replacing a failed node in a disaster recovery configuration

If you lose one of the appliances in a disaster recovery configuration, you can replace the appliance and restore the disaster recovery configuration by following this procedure.

About this task

If a disaster occurs such that the appliance in the main site is beyond repair, you delete the disaster recovery configuration while your queue manager runs on the recovery appliance. You then replace the appliance and restore the disaster recover configuration.

Procedure

Following the loss of the queue manager on the main site, take the following steps:

1. On the recovery appliance, run the following command:
   `makedrprimary -m QMname`

   Where *QMname* is the name of the queue manager.
2. Delete the disaster recovery configuration:
   `dltdrprimary -m QMname`
3. Run the following command to start the queue manager:
   `strmqm QMname`
4. Ensure that your applications reconnect to the queue manager on the recovery appliance. Provided that you have defined your channels with a list of alternative connection names, specifying your primary and secondary queue managers, then your applications will automatically connect to the new primary queue manager.
5. Replace the failed appliance on your main site and configure it for disaster recovery, see "Configuring the hardware for disaster recovery" on page 144.
6. Stop the queue manager:
   `endmqm QMName`
7. On the recovery site, make the queue manager the primary in a new disaster recovery configuration:
   `crtdrprimary -m QMName -r RecoveryName -i RecoveryIP -p Port`

   -m *QMName*
   Specifies the queue manager. The queue manager must be stopped when you run the command.

   -r *RecoveryName*
   Specifies the name of the replacement IBM MQ Appliance that you have installed on the main site.

   -i *RecoveryIP*
   Specifies the IP address of the replacement appliance.

   -p *port*
   Specifies the port that the data replication listener on both appliances uses.

   On successful completion, the command outputs the **crtdrsecondary** command

8. On the replacement appliance, run the command that was output by the `crtdrprimary` command on its successful completion, for example:
   ```
   crtdrsecondary –m QM1 –s 65536 –l myliveappl –i 198.51.100.24 –p 2015
   ```
   Synchronization of data from the recovery appliance to the main appliance begins.
9. On the recovery appliance, ensure that the queue manager is not running, and then make it the secondary queue manager:
   ```
   mkdrsecondary -m QMName
   ```
10. On the replacement appliance, make the queue manager the primary queue manager:
    ```
    mkdrprimary -m QMName
    ```
11. On the replacement appliance, start the queue manager:
    ```
    strmqm QMname
    ```
    You have now restored the configuration as it was before the failure at your main site.

    This procedure is shown in diagrammatic form in the following figure:

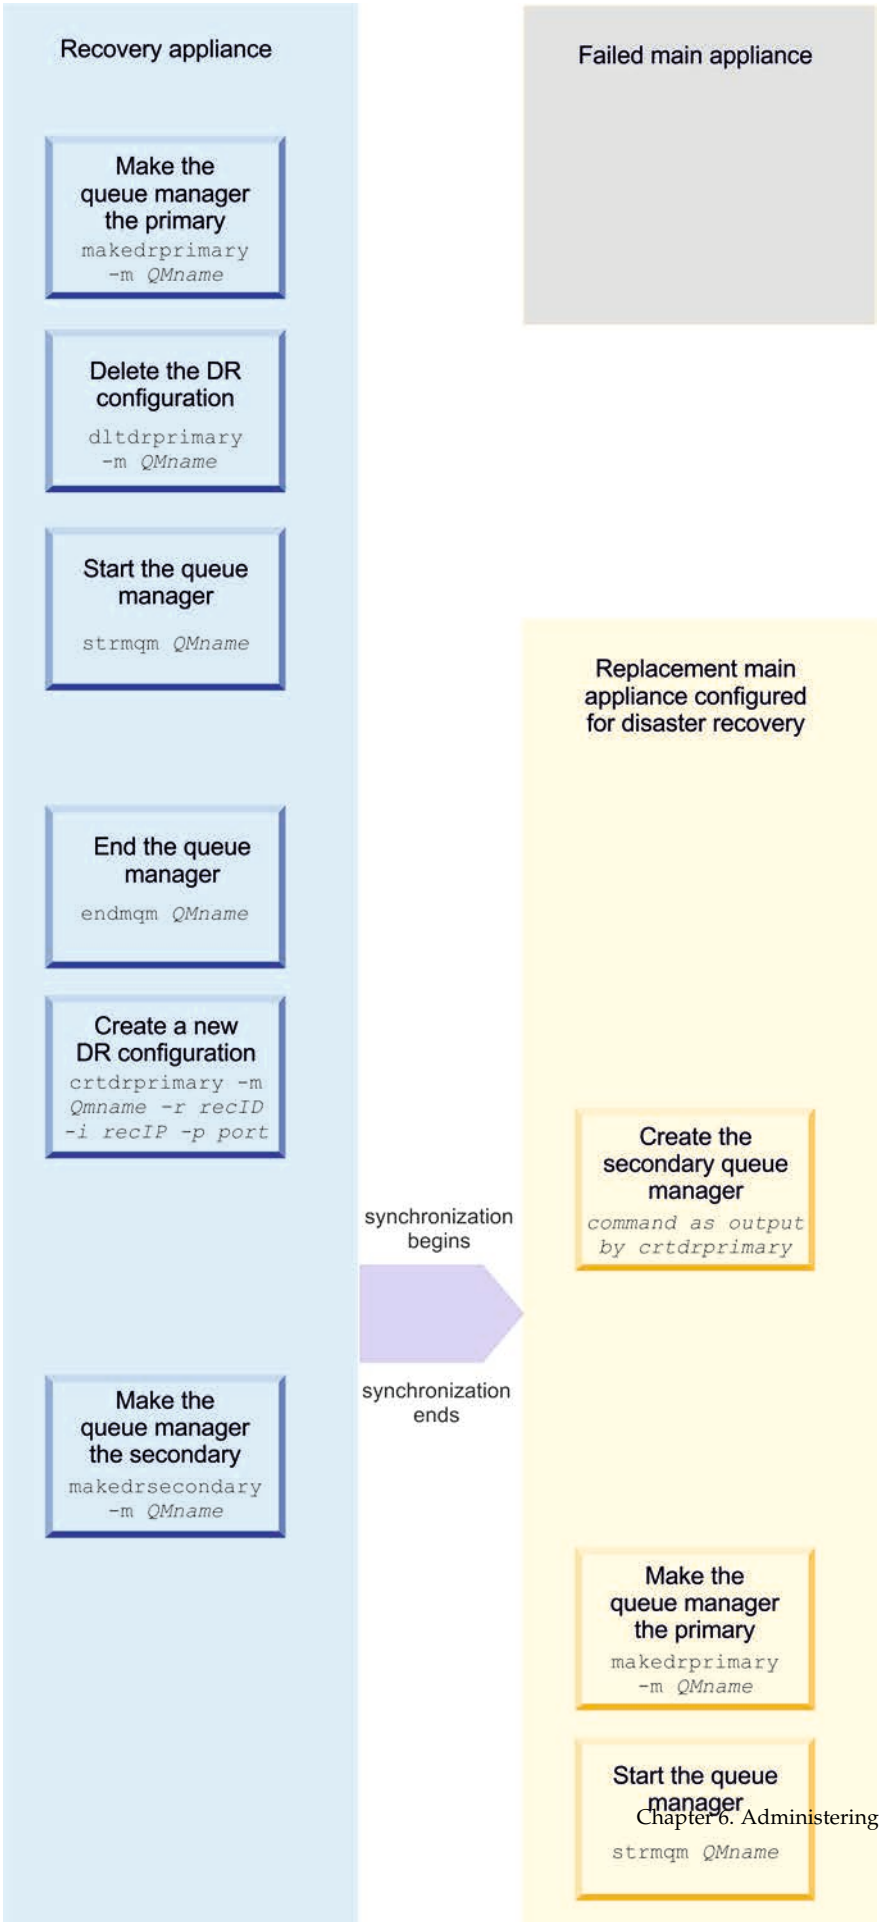
Recovery appliance
Make the queue manager the primary
makedrprimary -m QMname
Delete the DR configuration
dltdrprimary -m QMname
Start the queue manager
strmqm QMname
End the queue manager
endmqm QMname
Create a new DR configuration
crtdrprimary -m Qmname -r recID -i recIP -p port
Make the queue manager the secondary
makedrsecondary -m QMname
Failed main appliance
Replacement main appliance configured for disaster recovery
Create the secondary queue manager
command as output by crtdrprimary
synchronization begins
synchronization ends
Make the queue manager the primary
makedrprimary -m QMname
Start the queue manager
strmqm QMname

Testing the recovery appliance

You can test that the recovery appliance in a disaster recovery configuration is operating correctly without disrupting the main site.

About this task

You test the recovery appliance by disabling the replication interface between main and recovery appliances. You make the secondary queue manager into the primary and remove it from the disaster recovery configuration. You can then test the stand-alone queue manager. After testing is complete, you restore the replication interface and delete the queue manager. You then re-create the queue manager as the secondary queue manager in the disaster recovery configuration.

Procedure

1. Disable the replication link on one or both appliances:

   ```
   # ethernet eth20
   Modify Ethernet Interface configuration
   # admin-state disable
   # exit
   ```

   You can now work on the recovery appliance without affecting the main appliance.
2. On the recovery appliance, make the queue manager the primary:

   ```
   makedrprimary -m QMname
   ```

   Where *QMname* is the name of the queue manager.
3. Remove the queue manager from the disaster recovery configuration:

   ```
   dltdrprimary -m QMname
   ```
4. Start the queue manager:

   ```
   strmqm QMname
   ```
5. Connect applications to the queue manager and test that they work as expected.
6. End the queue manager:

   ```
   endmqm QMname
   ```
7. Delete the queue manager:

   ```
   dltmqm QMname
   ```
8. Restore the replication link between the main and recovery appliances:

   ```
   # ethernet eth20
   Modify Ethernet Interface configuration
   # admin-state enable
   # exit
   ```
9. Rerun the **crtsecondary** command that you used to create the secondary queue manager when you first created the disaster recovery configuration. The primary queue manager on the main appliance synchronizes its data with the secondary queue manager to bring it up to date.

Reversing disaster recovery roles

You can, if required, reverse the roles of the main and recovery appliances in your disaster recovery configuration.

About this task

When you reverse the roles of your main and recovery appliances, you convert primary queue managers into secondary queue managers on the original main appliance. You then convert secondary queue managers into primary queue managers on the original secondary appliance.

You should check the disaster recovery status of a queue manager before issuing the **makedrprimary** command, as the effects of the **makedrprimary** command depend on the status of the queue manager. Use the **status** *qmgrname* command to check status.

- If you issue **makedrprimary** for a queue manager that was Primary on an appliance that failed, or is Primary when the network between the appliances failed, then you will have to resolve a partitioned state if the appliance or network is restored.
- If you issue **makedrprimary** for a queue manager in the Inconsistent state, then the state of the queue manager will be reverted to the state when the previous Synchronization was started. This is intended to be used when the original appliance on which the queue manager was in the Primary role has failed and will not be restored.

Procedure

1. End the primary queue manager on the main appliance.

   ```
   endmqm QMname
   ```

   Where *QMname* is the name of the queue manager.
2. Convert the queue manager into a secondary queue manager:

   ```
   makedrsecondary -m QMname
   ```
3. On the secondary appliance, make the queue manager the primary queue manager:

   ```
   makedrprimary -m QMname
   ```

What to do next

Repeat these steps for each queue manager in your disaster recovery configuration.

Viewing the status of a disaster recovery queue manager

You can view the status of a queue manager in a disaster recovery configuration by using the **status** command on the command line, or by using the IBM MQ Console.

About this task

The **status** command returns information about the operational status of a specified queue manager in the disaster recovery configuration. The status can include the following information:

- The disaster recovery role of the queue manager (reported as `Primary` or `Secondary`).

- The current disaster recovery status:

  Normal
  The appliances in the disaster recovery configuration are operating normally.

  Synchronization in progress
  This status can mean that initial replication is completing, or there has been a failure of the disaster recovery replication network and the queue manager has switched into synchronization mode to catch up as quickly as possible.

  Partitioned
  Queue manager data on the appliances is out of step, and cannot be automatically resolved. The **makedrprimary** and **makedrsecondary** commands must be used to resolve the situation. When this status is displayed on one of the appliances in a disaster recovery pair, the other appliance might display the **remote appliance unavailable** status, because the connection was lost before it detected the partitioned status.

  Remote appliance(s) unavailable
  The status means that the connection to the other appliance in the disaster recovery configuration has been lost.

  Inactive
  The queue manager is in the secondary role on both appliances.

  Inconsistent
  This status is shown only when the queue manager is in the secondary role and synchronization is in progress. If you use the **makedrprimary** command on a queue manager that is in this state, the queue manager reverts to the snapshot of its data that was taken before it entered the inconsistent state.

  Reverting to snapshot
  This status is shown when the queue manager is in the secondary role, and the **makedrprimary** command is issued when the queue manager is in the `inconsistent` state. The queue manager is reverted to the current snapshot of its data such that it can run.

  Remote appliance(s) not configured
  This status is shown when the **crtdrprimary** command has been run, to specify that a queue manager has the primary role, but no **crtdrsecondary** command has been run on the other appliance in the disaster recovery pair.
- The percentage complete of a synchronization operation. This information is shown only when the status is `Synchronization in progress`.
- The estimated time at which a synchronization will complete. This information is shown only when the status is `Synchronization in progress`.
- The amount of out-of-sync data that exists on this instance of the queue manager. This is the amount of data written to this instance of the queue manager since it entered the partitioned state. This information is shown only when the status is `Partitioned`.
- The percentage complete of a reversion to snapshot operation. This information is shown only when the status is `Reverting to snapshot`.

Procedure

- To view the DR status of a queue manager by using the command line interface:
  1. Enter the IBM MQ administration mode by entering the following command:
     `mqcli`
  2. View the status of a DR queue manager by entering the following command from one of the appliances:
     `status` *QMgrName*
     Where:
     QMgrName
     Specifies the name of the DR queue manager that you want to view the status of.
  3. Exit the IBM MQ administration mode by entering the following command:
     `exit`
- To view the DR status of a queue manager by using the IBM MQ Console:
  1. Open the console and find the widget that displays the queue manager.
  2. Select the queue manager in the widget and select the properties icon from the tool bar .
  3. In the properties window, click the **Disaster recovery status** section to open it.

Managing files by using the IBM MQ Appliance web UI

You can use the file management interface of the IBM MQ Appliance web UI to view, transfer, and edit files on the appliance.

Accessing the file manager

To access the file manager, start the IBM MQ Appliance web UI and click the **File Management** tab.

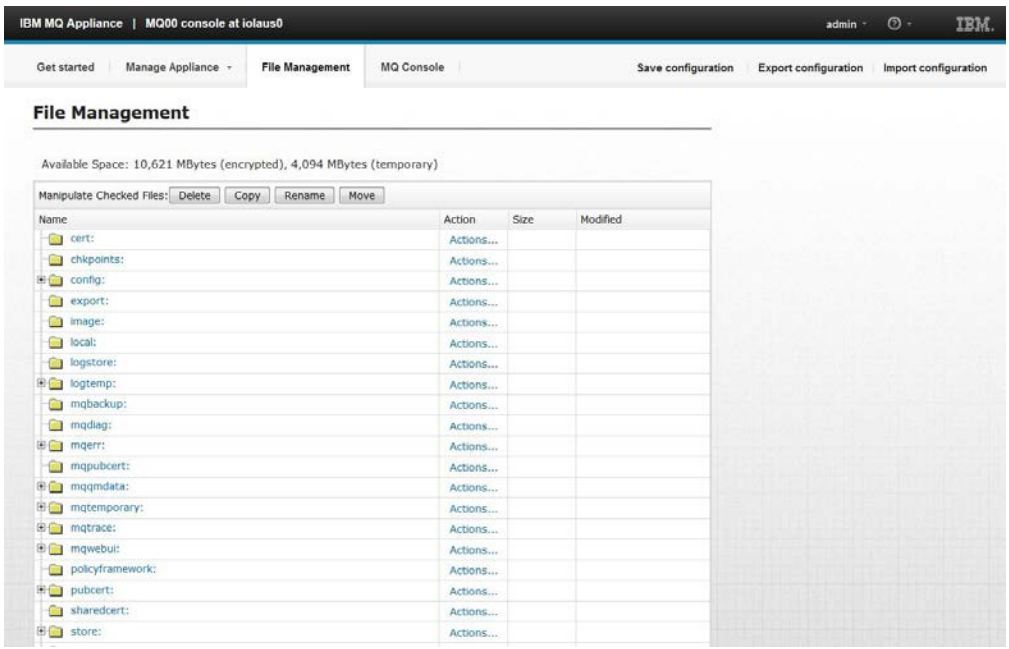

Displaying directory contents

To show the contents of a directory, click the directory name in the file management interface.

Uploading files from your workstation to the appliance

To upload files from your workstation to the appliance:

1. Click **Actions** in the row of the directory that you want to copy a file to.
2. Select **Upload Files**.
3. Define the file to upload to the appliance:
   a. Click **Browse** to locate the file that you want to upload.
   b. Optionally modify the file name of the destination file by typing a new name in the **Save as** field.
   c. Click **Add**.
4. To copy additional files, repeat the previous step.
5. Optionally select **Overwrite Existing Files**.
6. Click **Upload**.

Retrieving files from a remote location

To retrieve a file from a remote location and upload it to the appliance:

1. Click **Actions** in the row of the directory that you want to copy a file to.
2. Select **Fetch Files**.
3. Specify the location of the file in the **Source URL** field.
4. Specify the target file name in the **Save as** field.
5. Optionally select **Overwrite Existing Files**.
6. Click **Fetch**.

Copying files

To copy files from one directory to another directory on the appliance:

1. Locate the directory that contains the files to be copied.
2. Select the check box next to the file name.
3. Scroll to the top or bottom of the panel and click **Copy**.
4. From the **New Directory Name** list, select the target directory.
5. In the **New File Name** field, enter the target file name, if different.
6. Optionally select **Overwrite Existing Files**.
7. Click **Confirm Copy**.

Renaming files

To rename a file on the appliance:

1. Locate the directory that contains the file to be renamed.
2. Select the check box next to the file name.
3. Scroll to the top or bottom of the panel and click **Rename**.
4. In the **New File Name** field, enter the target file name.
5. Optionally select **Overwrite Existing Files**.
6. Click **Confirm Rename**.

Moving files

To move files from one directory to another directory on the appliance:

1. Locate the directory that contains the files to be moved.
2. Select the check box next to the file name.
3. Scroll to the top or bottom of the panel and click **Move**.
4. From the **New Directory Name** list, select the target directory.
5. Optionally select **Overwrite Existing Files**.
6. Click **Confirm Move**.

Deleting files

To delete files on the appliance:

1. Locate the directory that contains the files to be deleted.
2. Select the check box next to the file name.
3. Scroll to the top or bottom of the panel and click **Delete**.
4. Click **Confirm Delete**.

Viewing files

To view a file on the appliance:

1. Locate the directory that contains the file.
2. Click the file name to view its contents in your browser.

Editing files

You can edit some, but not all, of the files on the appliance. To edit a file:

1. Locate the directory that contains the file.

2. If the file can be edited, there is an edit link on the same row as the file name. Click **Edit**.
3. The file opens in preview mode. Click **Edit** to edit the file.
4. Edit the file as required.
5. Click **Submit**.
6. Click **Close**.

Downloading files from the appliance

To download a file from the appliance to your workstation, you use the download controls in your browser. So, for example, in Firefox you right-click a file name, select **Save link as** and browse for a location on your workstation to save the link to.

Watchdog timer

The IBM MQ Appliance has a baseboard management controller (BMC) that provides a watchdog timer.

The watchdog timer allows you to detect and recover from a serious malfunction on the appliance, even if the appliance is at a remote location. When the appliance is running normally, the appliance firmware informs the BMC that all is well every few seconds. If the BMC receives no such notification for a specified time (by default, twenty minutes), it restarts the appliance.

If you want to change the default behavior of the watchdog timer, or implement some of the other available features, you can configure the BMC.

You use the Intelligent Platform Management Interface (IPMI) to configure the BMC. Commands that are sent over IPMI are independent of the appliance CPU, firmware, and operating system. The BMC can still be accessed when the appliance is powered off (provided that it is plugged into power).

You must meet the following requirements before you can configure the BMC:

- You must use the mgt0 interface on the appliance for your IPMI connection, see "IPMI LAN channel commands" on page 538.
- You must create a special IPMI user, "IPMI user commands" on page 536.
- You require a remote system, for example a Linux host, running a suitable IPMI client. (The examples use a Linux command line IPMI client called ipmitool, see http://linux.die.net/man/1/ipmitool).

Examples

The following examples show basic watchdog timer configuration, by using ipmitool commands.

The following command queries the state of the watchdog timer:

```
ipmitool -L operator -I lanplus -H ipmi_channel_IP -U ipmi_user
-P ipmi_password mc watchdog get
```

Where:

- *ipmi_channel_IP* is the IP address that you allocated to the appliance when you configured the IPMI interface on mgt0.
- *ipmi_user* is the name of the ipmi user that you configured on the appliance.

- *ipmi_password* is the password for the ipmi user.

The command returns information similar to the following example:

```
Watchdog Timer Use:          SMS/OS (0x44)
Watchdog Timer Is:           Started/Running
Watchdog Timer Actions:     Hard Reset (0x01)
Pre-timeout interval:        0 seconds
Timer Expiration Flags:     0x00
Initial Countdown:           1200 sec
Present Countdown:           1199 sec
```

Where:

Watchdog Timer Is
Reports the current running state of the watchdog timer.

Watchdog Timer Action
Describes what is done when the timer reaches 0. The default is to restart the appliance.

Initial Countdown
The total timer wait time.

Present Countdown
The current timer value.

The following command disables the watchdog timer:

```
ipmitool -L operator -I lanplus -H ipmi_channel_IP -U ipmi_user
-P ipmi_password mc watchdog off
```

You receive a message confirming that the watchdog timer is disabled:

```
Watchdog Timer Shutoff successful -- timer stopped
```

The following command reenables the timer by setting it to its initial state:

```
ipmitool -L operator -I lanplus -H ipmi_channel_IP -U ipmi_user
-P ipmi_password mc reset warm
```

You receive a message confirming the reset:

```
Sent warm reset command to MC
```

Chapter 7. Migrating and consolidating

You can consolidate your IBM MQ estate by migrating existing queue manager configurations onto the IBM MQ Appliance.

The IBM MQ Appliance is designed to be a good candidate for consolidation scenarios, where an existing diverse estate of IBM MQ queue managers and applications is converged in a messaging hub architecture. Features of the environment that make the appliance ideal for this use case include the system performance tuning for client connectivity, high availability tooling, and segmentation available by using fixed storage allocations for queue managers.

A number of factors need consideration when you plan such a migration/consolidation exercise, depending on your previous IBM MQ configuration. The steps that are described in the following topics must be tailored to the particular environment that is being consolidated or migrated.

Moving queue managers from other IBM MQ platforms

Consolidation of your IBM MQ estate means moving your queue managers from their various platforms to your IBM MQ Appliance.

IBM MQ Appliance V8.0 is compatible with IBM MQ V8.0. Follow these instructions if you are moving queue managers from an IBM MQ V8.0 platform.

You use the **dmpmqcfg** command on your source system to save the configuration of a queue manager. Running **dmpmqcfg** records a series of MQSC commands that you later run with the **runmqsc** command. For information about MQSC commands, see MQSC commands in the IBM MQ documentation. You create a new queue manager on your target appliance, and create a connection to it on your source system. You then use the **runmqsc** command on the source system to configure the remote queue manager.

As part of moving a queue manager, you must carefully check the details that you are exporting. If there are features in the export that are not supported on IBM MQ Appliance, you must take action to remedy this. In particular, note you cannot run applications or services on the appliance. You must move such functionality to a client application.

If you move queue managers that are part of a distributed configuration, you must update channel definitions on other queue managers in the configuration to point to the new location of the moved queue manager on the appliance.

The following topics contain detailed instructions for moving queue managers from different types of platform.

Note: These instructions assume that you are moving queue managers from platforms other than z/OS, but the general principles also apply to migrating from z/OS.

© Copyright IBM Corp. 2015, 2016

Moving a queue manager

Follow these instructions to move a queue manager from an IBM MQ V8.0 system to an IBM MQ Appliance.

About this task

You move a queue manager by re-creating it on the target system. The procedure re-creates the configuration of the queue manager, it does not attempt to re-create the current state of the queue manager by, for example, unloading and reloading queues.

In these instructions, the source system is the system that you are moving the queue manager from. The target system is the IBM MQ Appliance.

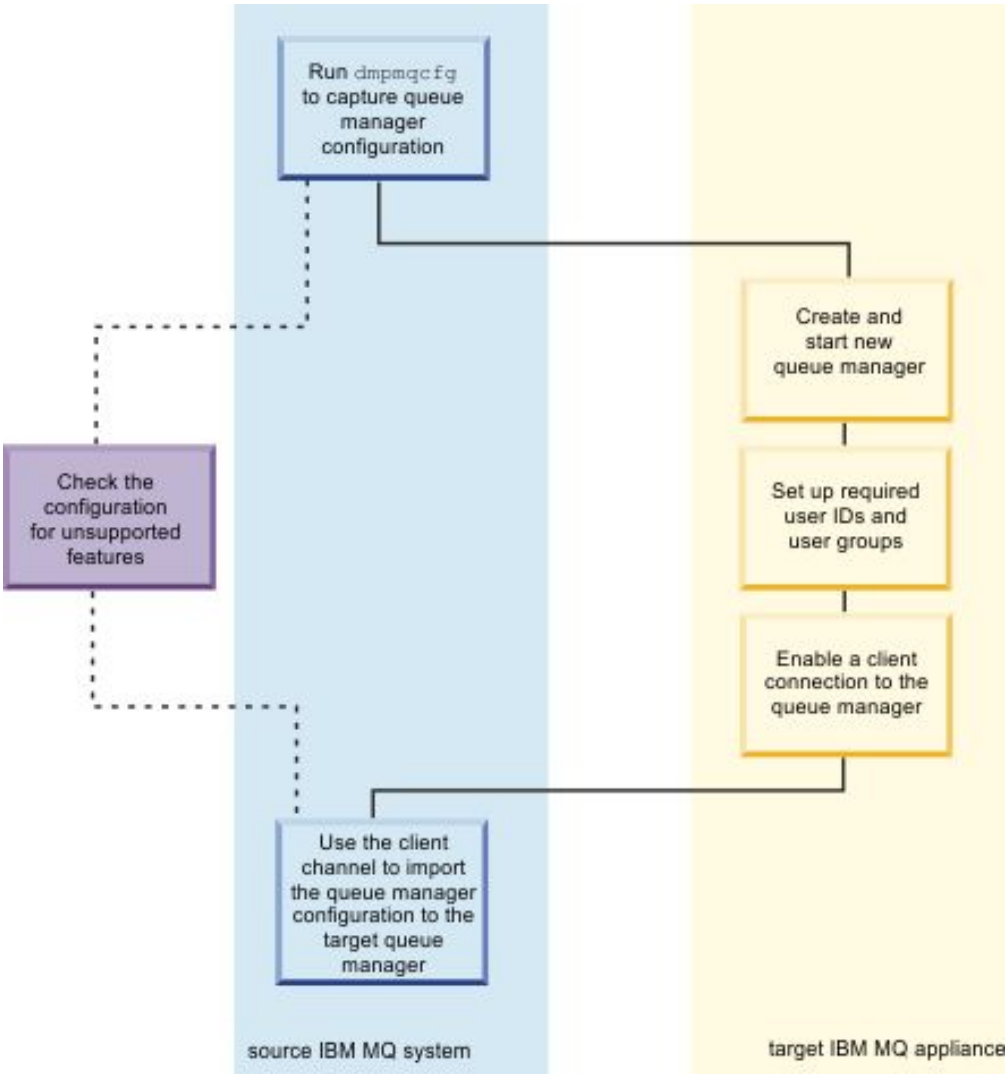

Procedure

1. Log in to the source system as a user in the IBM MQ administrators (mqm) group.
2. Save the configuration information of the queue manager that you want to move by typing the following command:

```
dmpmqcfg -a -m QM_name > QM_file
```

Where:

- *QM_name* is the name of the queue manager that you want to move.
- *QM_file* is the name and path of a local file on the source system that the configuration information is written to.

3. If the queue manager is part of a distributed configuration, quiesce the queue manager. Ensure that there are no messages in flight then stop the queue manager.
4. Create and start a new target queue manager on the IBM MQ Appliance. You can use the IBM MQ Console to do this action, see "Using the IBM MQ Console" on page 157, or you can use MQSC commands, with the required name and attribute values. If you want to use MQSC commands, you must complete the following steps:
   a. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
   b. Log in as a user in the administrators group.
   c. Type the following command to open the IBM MQ command line interface:

      ```
      mqcli
      ```
5. Set up any user IDs that are required by the queue manager that you are moving.
6. Enable a client connection to the target queue manager. You must define and start a TCP listener, define an SVRCONN channel, and allow administrator access to the queue manager by using this channel. You can use the IBM MQ Console to do this action, see "Using the IBM MQ Console" on page 157, or you can use MQSC commands. If the source IBM MQ system has IBM MQ Explorer, try using it to add a remote queue manager definition for the target queue to check that the client connection is working.
7. Ensure that your exported queue manager configuration is compatible with the target IBM MQ Appliance. Follow the process in "Handling incompatible features in the queue manager" on page 244. Edit the file that contains the queue manager configuration information if necessary.
8. Import the source queue manager configuration into the target queue manager. You run these steps on the source system:
   a. Define an environment variable that is named MQSERVER to identify the channel that connects to the target queue manager. For example, the value of MQSERVER could be set to:

      ```
      SYSTEM.ADMIN.SVRCONN/TCP/9.20.233.217(1414)
      ```
   b. Run the following command to replay on the target queue manager the commands that were exported from the source queue manager:

      ```
      runmqsc -c QM_name  < QM_file
      ```
9. Restore the attributes that were masked in the **dmpmqcfg** output and that you identified when you checked the output (as described in "Handling incompatible features in the queue manager" on page 244). You restore attributes by using the client connection from the source system. You can either use IBM MQ Explorer, or start **runmqsc** interactively in client mode, and then input MQSC commands:

   ```
   runmqsc -c QM_name
   ```
10. Stop and restart the queue manager on the target IBM MQ Appliance and ensure that it starts cleanly.

Moving queue managers secured by using TLS

You must take additional steps when you move queue managers that are secured by using TLS.

About this task

When you move a secure queue manager to IBM MQ Appliance, you must re-create the repository on the appliance and regenerate certificates and keys. The repository is created when you create the queue manager on the appliance; you must take steps to regenerate certificates and keys. You then redistribute those certificates and keys to the various queue managers and clients that want to communicate with each other.

The following procedure describes a scenario that requires certificate exchange using a self-signed certificate. If you are using certificates signed by a CA, you require extra steps to request a certificate and to import the signed certificate and any other certificates required to form the chain of trust.

Procedure

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
2. Log in as a user in the administrators group.
3. Type the following command to open the IBM MQ command line interface shell:

   ```
   mqcli
   ```
4. Type the following command to generate a self-signed certificate, and extract it:

   ```
   mqa(mqcli)# createcert -m qmname -label labelname -dn "CN=Issuer,OU=Certificate Authority,O=organ
   ```

   For example:

   ```
   mqa(mqcli)# createcert -m REGA -label ibmwebspheremqrega -dn "CN=Issuer,OU=Certificate Authority,
   ```
5. Type **exit** to exit the IBM MQ command line interface shell, and type the following command to open the appliance configuration shell:

   ```
   config
   ```
6. Copy the new certificate that you created in step 4 to any queue manager or client machines that need a TLS connection to the queue manager:

   ```
   mqa(config)# copy mqpubcert:///certificate_source scp://certificate_destination
   ```

   For example:

   ```
   mqa(config)# copy mqpubcert:///REGA_ibmwebspheremqrega scp://myuser@9.20.120.129//build/exported_
   ```
7. If the queue manager is part of a distributed configuration, copy certificates from remote queue managers to the appliance. Enter the following command:

   ```
   mqa(config)# copy scp:certificate_source mqpubcert:///certificate_destination
   ```

   For example:

   ```
   mqa(config)# copy scp://myuser@9.20.120.129//build/exported_certficates/ibmwebspheremqregb.p12 mq
   ```
8. Open the IBM MQ CLI shell again, and type the following command to add certificates that you copied in the previous step to the repository:

   ```
   mqa(mqcli)# addcert -m qmname -label qmlable -file remoteqm_certificate -format ascii
   ```

   For example:

   ```
   mqa(mqcli)# addcert -m REGA -label ibmwebspheremqregb -file ibmwebspheremqregb.p12 -format ascii
   ```

9. On each of the systems that need to make a TLS connection with the queue manager on the IBM MQ Appliance, issue the commands to delete the original certificate and add the new one copied from the appliance.

Planning for incompatible features in the queue manager

It is possible that not all features in your source queue manager are supported by the target IBM MQ Appliance. You should take time to plan how you intend to handle any incompatible features.

For help with planning how to handle incompatible features, consult the following topics:

- "Differences between administering an IBM MQ Appliance and an IBM MQ installation" on page 13
- "Moving a queue manager" on page 240
- "Using an IBM MQ client" on page 192

user IDs and groups

As part of moving the queue manager, you must identify any user IDs and groups that the queue manager configuration includes and re-create them on the IBM MQ Appliance. If different user IDs and groups are created on the appliance, then you must make the appropriate changes to the **dmpmqcfg** output.

Special considerations for moving a queue manager from z/OS

In most cases it is not appropriate to move a queue manager from z/OS to an appliance, because the connecting applications (for example in batch, CICS, IMS and Db2 environments) must be locally bound to a z/OS queue manager running on the same LPAR as the application.

Queue managers on z/OS are likely to have several z/OS-specific attributes that are not supported on IBM MQ Appliance. You must remove or comment out such attributes.

These changes do not ensure that the migrated queue manager is functionally equivalent to the original queue manager on z/OS. You must consider each of the attributes that are not supported by the new queue manager to decide whether its value is significant for your applications, and if the behavior of the object in the new queue manager, without this attribute, is acceptable. In some cases, it might be necessary to define different objects or to set other values to achieve the same effect. This consideration also applies to differences in the default value of some attributes. For example, queues on z/OS default to non-shared so there might be several statements that replace queues, including default system queues, with non-shared versions. This action might be the right thing to do if your applications rely on this characteristic, or it might be the wrong thing to do because you want to preserve the default behavior of the appliance queue manager.

Inspecting `qm.ini` file for the source queue manager

Examine the `qm.ini` file and make a note of any settings that cannot be made by running the commands in the **dmpmqcfg** output. These settings might include, for example, log file settings. Particularly note any exit information in the configuration. IBM MQ Appliance does not support exits, so this functionality must be substituted. For example, channel exits can be replaced by channel auth

records, and API exits might be replaced by activity trace. For more information about `qm.ini`, see Queue manager configuration files, qm.ini in the IBM MQ documentation.

Applications

Applications cannot be run on the IBM MQ Appliance. You must plan to migrate any applications that are local to the queue manager to a client system. Such applications need to be rebuilt so that they can connect to the queue manager from another machine by using client connections. If any applications are run as triggered processes, they must also be converted to run on a client machine. In that case, it is necessary to run the trigger monitor in client mode and to alter the queue manager's process definitions accordingly. For help, see runmqtmc and Managing objects for triggering in the IBM MQ documentation.

Exits and services

The IBM MQ Appliance does not support exits or services that are defined in the queue manager configuration. You must plan to migrate exits and services to equivalent functionality on a client system. For guidance, see "Moving a queue manager" on page 240 and "Using an IBM MQ client" on page 192.

Channels that use SSLv3 CipherSpecs

By default, IBM MQ v8.0.0.2 does not support SSLv3 and related CipherSpecs. See Deprecation: SSLv3 protocol. If you move a queue manager to the IBM MQ Appliance that has one or more channels that use SSLv3, you can take action to enable support for SSLv3. You can take one of the following actions:

- Set the environment variable AMQ_SSL_V3_ENABLE=1. See "Configuring environment variables" on page 124 for details of how to set environment variables on the appliance.
- Change the SSL stanza in the `qm.ini` file to re-enable the SSLv3 CipherSpecs, for example:

  ```
  mqa# mqcli
  mqa(mqcli)# setmqini -m QMGR -s SSL -k AllowSSLV3 -v YES
  ```

  See "Adding a value to the configuration file" on page 245 for details of using the **setmqini** command to edit the `qm.ini` file on the appliance.

Handling incompatible features in the queue manager

You must check that the queue manager that you are moving to the IBM MQ Appliance is compatible with the appliance.

The **dmpmqcfg** command that you run on your source platform produces a series of MQSC commands that you run to re-create the queue manager on the target IBM MQ Appliance. Certain features are incompatible with the appliance, and you must check the **dmpmqcfg** output, and amend it if necessary, to deal with incompatible features.

The output from the **dmpmqcfg** command contains lines that are commented out by the asterisk (*) character. Many of these values are read-only values that are set when the queue manager is created. They cannot be affected by the commands in the **dmpmqcfg** output.

You must also check the configuration file (`qm.ini`) for the source queue manager and make a note of any non-default attributes that cannot be set by the ALTER QMGR command, and so are not recorded in the output from **dmpmqcfg**.

Substitute appropriate values for masked values

The output from the **dmpmqcfg** command might include one or more masked values. If these values were replayed in commands, they would not correctly re-create the objects configuration. The values are masked to prevent sensitive data, such as passwords, from being included in clear text in the configuration dump.

Before you replay the configuration, first check the output for masked parameters such as SSLCRYP, PASSWORD, or LDAPPWD that are commented. You must use additional commands to substitute valid values.

Remove definitions of queue manager services

Queue manager services are not supported by IBM MQ Appliance. You must search the **dmpmqcfg** output for any DEFINE SERVICE or ALTER SERVICE commands and remove service definitions. Services can be replaced by code in client applications.

Remove changes to the CCSID

Remove any change to the queue manager CCSID in the ALTER QMGR command. The default CCSID for IBM MQ Appliance is 819. If you must change the CCSID, use a separate command and then restart the queue manager to ensure that all processes switch to the new CCSID.

Verify user IDs

Ensure that any user IDs specified in the commands are correctly defined on the IBM MQ Appliance. On Windows source systems, the user and group names might be in the form *name@domain*. This format is not supported on the appliance, so any such user IDs must be mapped to new user IDs on the appliance.

Remove changes to the SSLKEYR queue manager attribute

The SSLKEYR queue manager attribute is managed by the appliance, and should not be overwritten when you replay the commands to create the queue manager configuration.

Removing listeners from Windows queue managers

Where you are moving a queue manager from a source Windows system, you must remove any definitions for NETBIOS, SPX, and LU62 listeners from the **dmpmqcfg** output.

Editing qm.ini files

You cannot directly edit a queue manager `qm.ini` file on the IBM MQ Appliance. There are CLI commands, however, that you use to work with `qm.ini` files.

Adding a value to the configuration file

You can add a value or modify an existing value in the configuration file of a queue manager by using the **setmqini** command on the command line.

About this task

You can use the **setmqini** command to add values to the qm.ini file, which is used for general queue manager configuration, or the mqat.ini file, which is used to control application activity trace. The stanza that you specify as an argument to **setmqini** determines which file the value is written to.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   mqcli
2. Add or modify the value to the qm.ini file by entering the following command:
   setmqini -m *QMgrName* -s *Stanza* -k *KeyName* -v *Value*
   Where:

   QMgrName
   Specifies that the configuration file that is associated with the specified queue manager is to be modified.

   Stanza
   Specifies which stanza the value is to be added to.

   The following values for *Stanza* modify the qm.ini file:
   - Log
   - TCP
   - Channels
   - TuningParameters
   - SSL
   - Security
   - Subpool

   The following value for *Stanza* modifies the mqat.ini file.
   - AllActivityTrace

   Do not edit the qm.ini file to control the number of channels. Instead, use the MAXINST and MAXINSTC values on your SVRCONN channels. For more information, see "Queue manager configuration on the IBM MQ Appliance" on page 19.

   KeyName
   Specifies which key to add or modify.

   Ensure that the value of *KeyName* is correct before you use the command to add a key and value from the stanza. The value of *KeyName* is not validated. If incorrect values are specified in the qm.ini file, a subsequent attempt to start the queue manager might fail.

   See Configuring trace levels for details of keys that you can add or modify to the mqat.ini file.

   Value Specifies the value to add for the specified key name.

   If *Value* is a string that contains spaces, it must be enclosed in double quotation marks. Any double quotation marks that are used in the *Value* must be escaped by using a backslash (\ ).

Ensure that the value of *Value* is correct before you use the command to add a value to the stanza. The value of *Value* is not validated. If incorrect values are specified in the `qm.ini` file, a subsequent attempt to start the queue manager might fail.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the addition of the key `RemoteQueueAccessControl` with a value of `Xmitq` to the stanza `Security` in the `qm.ini` file of queue manager `QM1`:

```
setmqini -m QM1 -s Security -k RemoteQueueMangerAccessControl -v Xmitq
```

The following example shows the key `TraceLevel` being set to HIGH in the `mqat.ini` file.

```
setmqini -m QM1 -s AllActivityTrace -k TraceLevel -v HIGH
```

Deleting a value from a qm.ini file

You can delete a value from the `qm.ini` file of a queue manager by using the **setmqini** command on the command line.

About this task

You cannot delete an entire stanza from the `qm.ini` in a single command. To delete an entire stanza, you must delete each key individually from the stanza.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Delete the value from the `qm.ini` file by entering the following command:

   ```
   setmqini -m QMgrName -s Stanza -k KeyName -d
   ```

   Where:

   QMgrName
   Specifies that the `qm.ini` file that is associated with the specified queue manager is to be modified.

   Stanza
   Specifies which stanza the value is to be removed from.

   Valid values for *Stanza* are the following values:

   - Log
   - TCP
   - Channels
   - TuningParameters
   - SSL
   - Security
   - Subpool

   KeyName
   Specifies which key and associated value to delete.

Ensure that the value of *KeyName* is correct before you use the command to remove a key and value from the stanza. The value of *KeyName* is not validated. If you remove a value that is required by the queue manager, a subsequent attempt to start the queue manager might fail.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the deletion of the key name and associated value of `RemoteQueueAccessControl` in the `Security` stanza of the `qm.ini` file of queue manager `QM1`:

```
setmqini -m QM1 -s Security -k RemoteQueueAccessControl -d
```

Viewing a configuration file

You can view the contents of a single stanza or key in a queue manager configuration file by using the **dspmqini** command on the command line.

About this task

You can use the **dspmqini** command to view stanzas in the `qm.ini` file, which is used for general queue manager configuration, or the `mqat.ini` file, which is used to control application activity trace. The stanza that you specify as an argument to **dspmqini** determines which file you view.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. View the contents of the configuration file by entering one of the following commands:
   - To view the contents of the entire `qm.ini` file, enter the following command:

     ```
     dspmqini -m QMgrName
     ```

   - To view the contents of a single stanza of a `qm.ini` file, or of the `mqat.ini`file, enter the following command:

     ```
     dspmqini -m QMgrName -s Stanza
     ```

   - To view the contents of a single key of the `qm.ini` or `mqat.ini` file, enter the following command:

     ```
     dspmqini -m QMgrName -s Stanza -k KeyName
     ```

   Where:

   QMgrName
   Specifies that the file that is associated with the specified queue manager is to be viewed.

   Stanza
   Specifies the stanza that you want to view.

   The following values are valid for viewing stanzas in the `qm.ini` file:
   - Log
   - TCP
   - Channels

- TuningParameters
- SSL
- Security
- Subpool

The following value is valid for viewing the `mqat.ini` file:

- AllActivityTrace

KeyName
Specifies the name of the key that you want to view.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

- The following example shows viewing the `Channels` stanza in the `qm.ini` file for queue manager `QM1`:

  ```
  dspmqini -m QM1 -s Channels
  ```

- The following example shows viewing the value of the key name `ClusterQueueAccessControl` in the `Security` stanza of the `qm.ini` file of queue manager `QM1`:

  ```
  dspmqini -m QM1 -s Security -k ClusterQueueAccessControl
  ```

- The following example shows viewing the value of the key name `ActivityCount` in the `AllActivityTrace` stanza of the `mqat.ini` file of queue manager `QM1`:

  ```
  dspmqini -m QM1 -s AllActivityTrace -k ActivityCount
  ```

Transferring queue managers to other IBM MQ Appliances

You can transfer queue managers and associated data from one appliance to another. You can use the High Availability or Disaster Recovery features to assist in such transfers, which simplifies the procedure.

Transfer from an existing single appliance to a new single appliance

Follow this procedure to transfer queue managers from an existing IBM MQ Appliance to a new IBM MQ Appliance.

The following diagram gives an overview of the procedure.

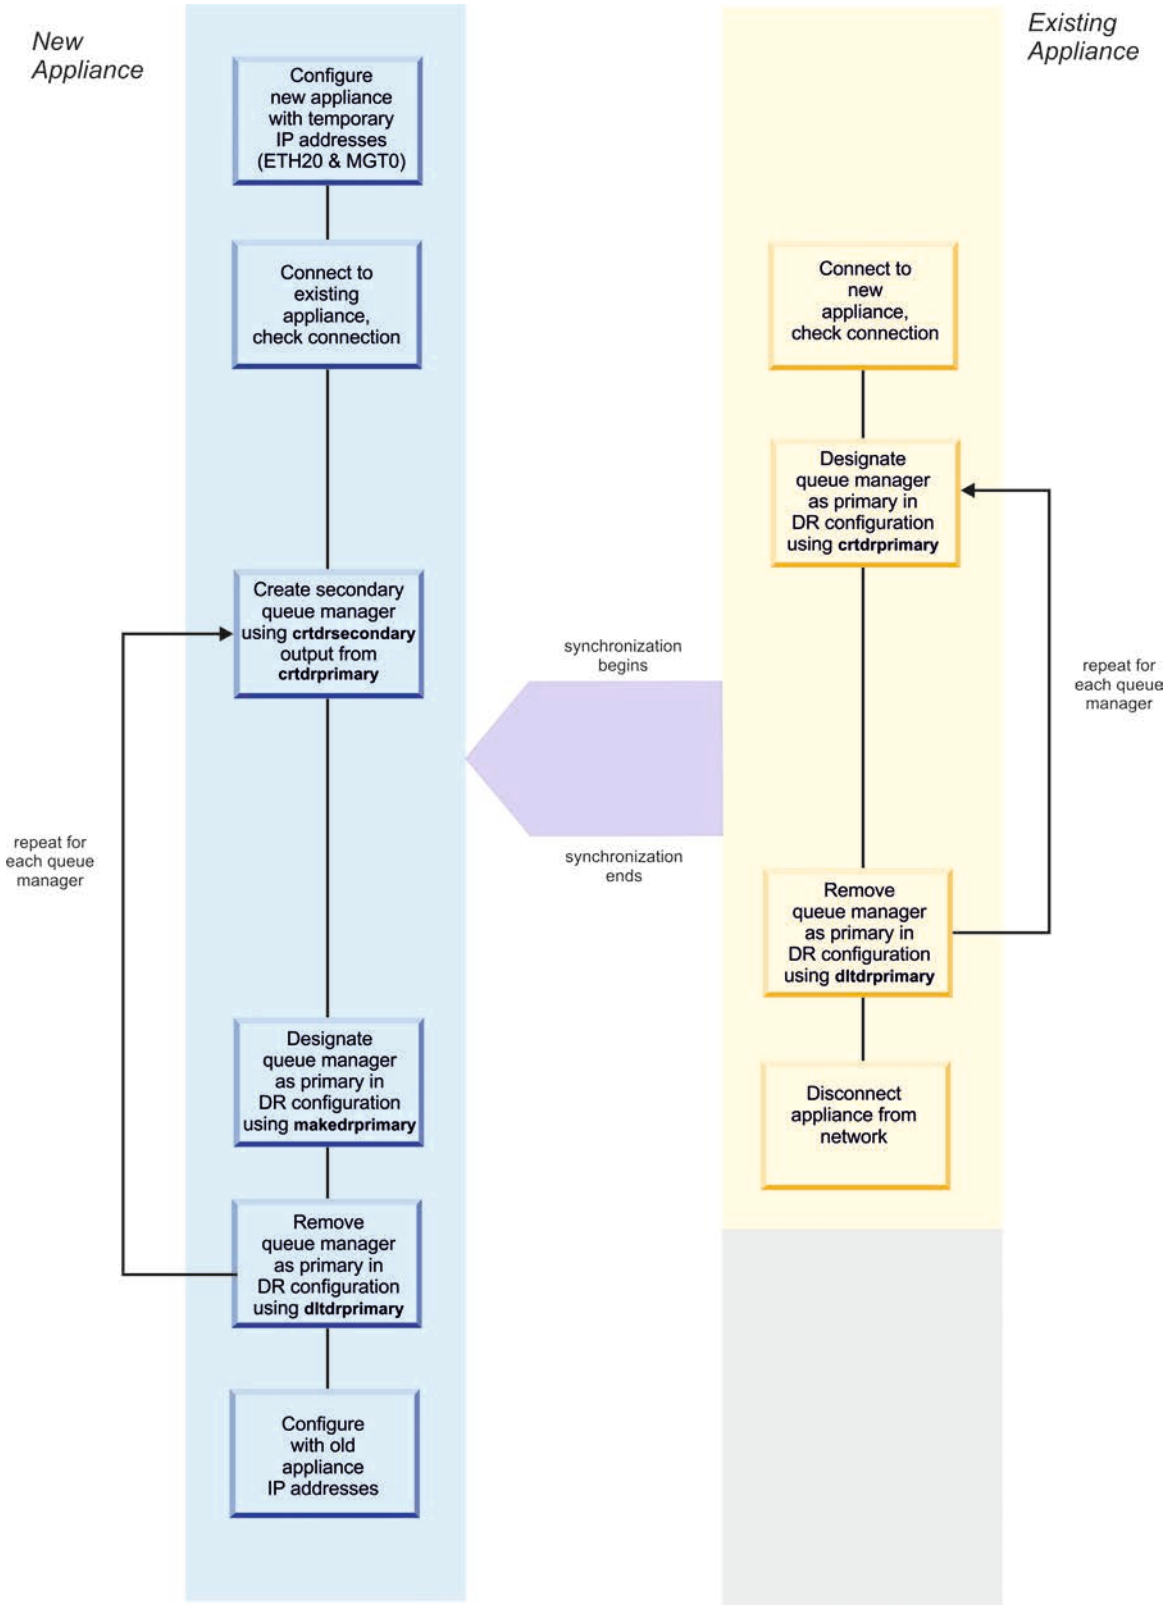

Figure 46. Transferring queue managers on single appliances

Preparing the new appliance

Follow these steps to prepare the new appliance for transfer of an IBM MQ queue manager and associated data from the old appliance.

About this task

To prepare the new appliance, you connect it to the old appliance and give the new appliance a temporary configuration for ports ETH20 and MGT0. When you have transferred all your queue managers, the new appliance will be reconfigured with the old appliance port details, and the old appliance is then retired.

Procedure

1. Power up the new appliance and attach the supplied serial console cable to the console port.
2. Configure the ETH20 port. This is the port over which you will migrate the queue manager data. If you are going to directly connect the two appliances, configure a temporary IP address and leave the gateway unconfigured. If you are connect to a 10 Gb network, specify an IP (CIDR) and a gateway for the connection. See “Configuring Ethernet interfaces by using the command line” on page 105 for help with configuring the Ethernet ports.
3. Configure the MGT0 port. Connect the port to a 1 Gb network, specifying an IP (CIDR) and a gateway for the connection.
4. Connect the ETH20 port on the new appliance to the ETH20 port on the old appliance using the supplied cable. (If the two appliances are not located near each other, ensure both are connected to a 10 Gb network.)
5. Check that you can communicate with the new appliance by pinging the IP address that was assigned to ETH20 on the existing appliance.
6. Ensure the MGT0 port on both appliances are connected to a 1 Gb network (this connection enables you to send commands to the appliance from a remote workstation).

Preparing the existing appliance

Follow these steps to prepare your existing appliance for transfer of an IBM MQ queue manager and associated data to the new appliance.

About this task

To prepare the old appliance, you connect it to the new appliance and ensure that the two appliances can communicate.

Procedure

1. If it is not already configured, configure the ETH20 port. This is the port over which you will migrate the queue manager data. If you are going to directly connect the two appliances, configure a an IP address and leave the gateway unconfigured. If you are connect to a 10 Gb network, specify an IP (CIDR) and a gateway for the connection. See “Configuring Ethernet interfaces by using the command line” on page 105 for help with configuring the Ethernet ports.
2. Ensure that the ETH20 port on the existing appliance is connected to the ETH20 port on the new appliance using the supplied cable. (If the two appliances are not located near each other, ensure both are connected to a 10 Gb network.)
3. Check that you can communicate with the new appliance by pinging the IP address that was assigned to ETH20 on the new appliance.

4. Ensure the MGT0 port on both appliances are connected to a 1 Gb network (this connection enables you to send commands to the appliance from a remote workstation).

Transferring a queue manager

Follow these steps to transfer a queue manager and associated data from one appliance to another appliance.

About this task

You use the commands that are normally used to set up a disaster recovery solution to transfer data between the two appliances.

On the existing appliance, you use a command that specifies that the queue manager to be transferred is the primary version of the queue manager. When you run this command, it outputs another command that you run on the new appliance to create a new, secondary version of the queue manager. After you run the command to create a secondary version of the queue manager, synchronization begins and transfers the queue manager data across the ETH20 link.

You repeat this procedure for each queue manager on your existing appliance.

Procedure

1. On the existing appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Stop the queue manager that you want to transfer:
      `endmqm queue_manager`
   c. Enter the following command:
      `crtdrprimary -m queue_manager -r new_appliance_name -i new_appliance_IP -p port`
      On successful completion, the command outputs a **crtdrsecondary** command.
   d. Restart the queue manager:
      `strmqm queue_manager`
2. On the new appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Enter the **crtdrsecondary** command, exactly as output in step 1c. Synchronization of data from the old appliance to the new appliance begins.
3. On either appliance, check the progress of the synchronization by using the **status** command.
   `status queue_manager`

Post transfer tasks

After you have transferred all of your queue managers from your existing appliance to your new appliance, you retire your existing appliance and configure the new appliance to take its place.

About this task

Your first step is to use more disaster recovery commands to remove the pairing between the two machines. You then disconnect your old appliance from the network. Finally, you reconfigure the new appliance so that it uses the same IP addresses as the old appliance.

Procedure

1. On the old appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Stop the queue manager that you want to work with:
      `endmqm queue_manager`
   c. Enter the following command to specify that a queue manager is no longer the primary instance of that queue manager and remove it from the temporary disaster recovery configuration:
      `dltdrprimary -m queue_manager`
   d. Repeat these steps for every queue manager on your appliance.
2. On the new appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Enter the following command to specify that a queue manager is now the primary instance of that queue manager:
      `makedrprimary -m queue_manager`
   c. Stop the queue manager:
      `endmqm queue_manager`
   d. Enter the following command to remove the queue manager from the temporary disaster recovery configuration:
      `dltdrprimary -m queue_manager`
   e. Repeat these steps for every queue manager on your appliance.
3. Disconnect your old appliance from any networks that it is connected to.
4. Reconfigure your new appliance so that it uses the IP addresses previously used by the old appliance. This step ensures that any IP address used in an IBM MQ channel now connects to the new appliance. See “Configuring Ethernet interfaces by using the command line” on page 105 for help with reconfiguring your new appliance.

Transfer from an existing appliance in a disaster recovery configuration

Follow this procedure to transfer queue managers from an existing IBM MQ Appliance to a new IBM MQ Appliance that is part of a disaster recovery configuration.

The following diagram gives an overview of the procedure.

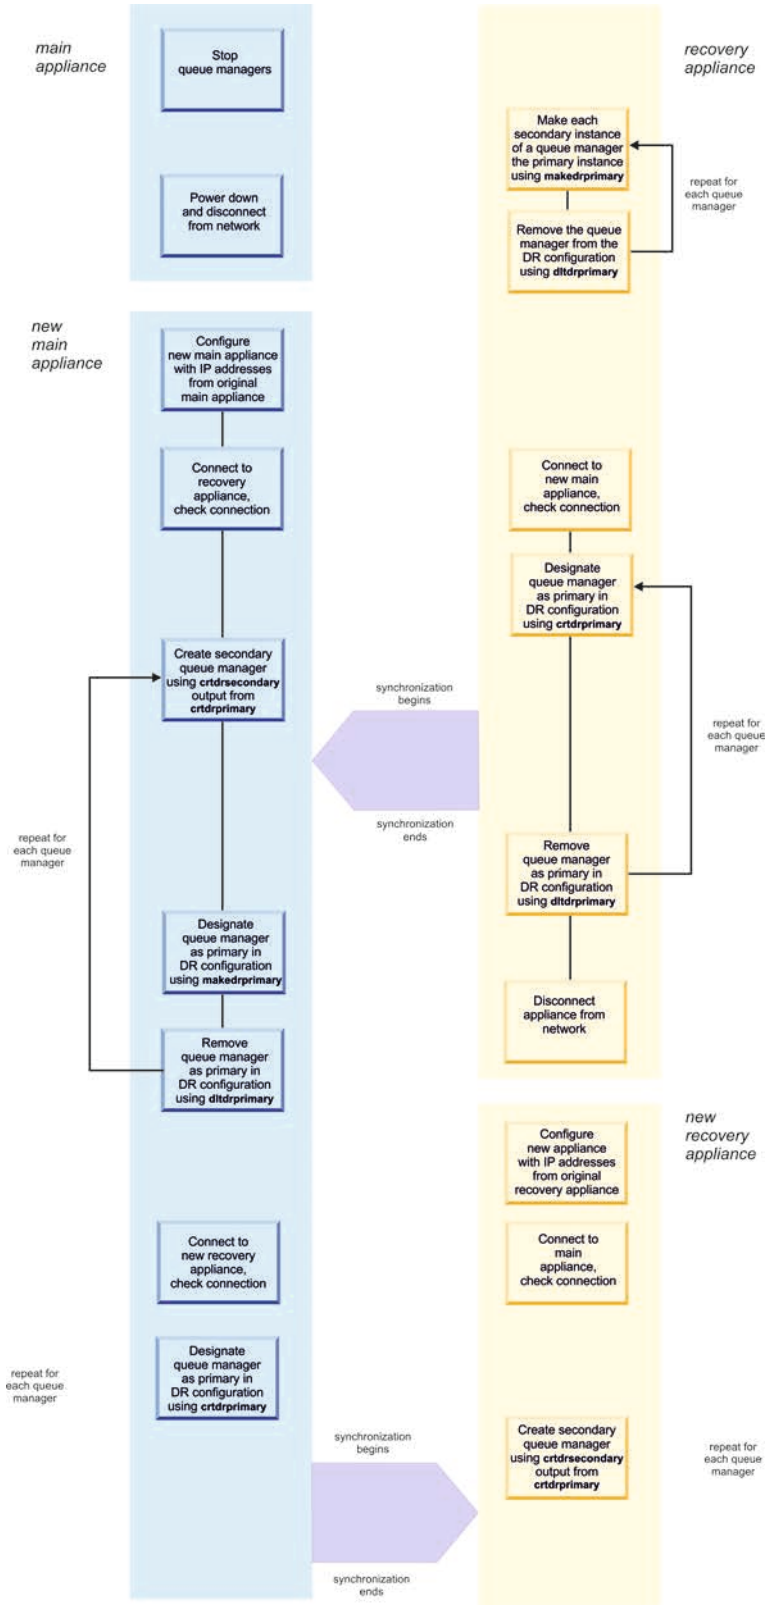

Figure 47. Transferring from an existing appliance in a disaster recovery configuration

Transferring from the main appliance

Follow this procedure to transfer queue managers and associated data from your existing main appliance.

About this task

This procedure uses the following terminology:

- Main appliance - the appliance located in your main data center that runs the primary instances of the queue managers.
- Recovery appliance - the appliance in your back up data center that has secondary instances of the queue managers.

To prepare your main appliance to transfer queue managers to a new, replacement appliance, you stop all the queue managers on the appliance, and run the secondary instances as primaries on your recovery appliance. You then delete the disaster recovery pairing between the two appliances.

The main appliance now has no active queue managers, you can disconnect it and power it down.

Procedure

1. On the main appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Stop the queue manager that you want to work with:
      `endmqm` *`queue_manager`*
2. On your recovery appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Make the queue manager the primary instance:
      `makedrprimary` *`queue_manager`*
   c. Start the queue manager:
      `strmqm` *`queue_manager`*
   d. Remove the queue manager from the disaster recovery configuration:
      `dltdrprimary -m` *`queue_manager`*
   e. Repeat these steps for each queue manager.
3. After you have transferred all your queue managers from the main appliance, you can disconnect it from the network and power it down.

Preparing the new main appliance

Follow these steps to prepare the new appliance for transfer of an IBM MQ queue manager and associated data from the recovery appliance.

About this task

To prepare the new appliance you power up the appliance and configure the Ethernet ports as they were configured for the original main appliance.

Procedure

1. Power up the new main appliance and attach the supplied serial console cable to the console port.
2. Configure the ETH20 port. This is the port over which you will migrate the queue manager data. Configure it with the same details that the appliance you

are replacing was configured. See "Configuring Ethernet interfaces by using the command line" on page 105 for help with configuring the Ethernet ports.

3. Configure the MGT0 port. Connect the port to a 1 Gb network, and configure it with the same details that the appliance you are replacing was configured. (This connection enables you to send commands to the appliance from a remote workstation.)
4. Configure the remaining Ethernet ports as they were originally configured for the main appliance.
5. Connect the ETH20 port to the 10 Gb network that the ETH20 port of the recovery appliance is connected to.
6. Connect the MGT0 port to the 1 Gb network that the MGT0 port of the recovery port is connect to.
7. Ping the secondary appliance to check the connections.
8. Make other network connections as required to replicate the configuration of the original main appliance.

Transferring queue managers to the new main appliance

Follow these steps to transfer a queue manager and associated data from the recovery appliance to your new main appliance.

About this task

After following the procedure to this point, your queue managers are currently running on your recovery appliance, and are not part of a disaster recovery configuration. You now use a command for each queue manager to designate that it is the primary instances on the recovery appliance. A command is output at the successful completion of this command that you run on the new main appliance to create a secondary instance of the queue manager there and synchronize it with the primary instance across the ETH20 link.

You repeat this procedure for each queue manager on your recovery appliance.

Procedure

1. On the recovery appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Stop the queue manager that you want to transfer:

      ```
      endmqm queue_manager
      ```

   c. Enter the following command:

      ```
      crtdrprimary -m queue_manager -r new_appliance_name -i new_appliance_IP
      -p port
      ```

      On successful completion, the command outputs a **crtdrsecondary** command.
   d. Restart the queue manager:

      ```
      strmqm queue_manager
      ```

2. On the new mains appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Enter the **crtdrsecondary** command, exactly as output in step 1c. Synchronization of data from the old appliance to the new appliance begins.
3. On either appliance, check the progress of the synchronization by using the **status** command.

   ```
   status queue_manager
   ```

4. On the secondary appliance, stop each queue manager and enter the following command to specify that the queue manager is no longer the primary instance and remove it from the disaster recovery configuration:

   ```
   dltdrprimary -m queue_manager
   ```

5. On the new main appliance enter the following command for each queue manager to specify that it is now the primary instance of the queue manager:

   ```
   makedrprimary -m queue_manager
   ```

6. Stop each queue manager:

   ```
   endmqm queue_manager
   ```

7. Enter the following command for each queue manager to remove it from the disaster recovery configuration:

   ```
   dltdrprimary -m queue_manager
   ```

Replacing the recovery appliance

Follow these steps to prepare the new recovery appliance to take its place in the disaster recovery configuration.

About this task

To prepare the new recovery appliance you power up the new appliance and configure the Ethernet ports as they were configured for the original recovery appliance.

You then work on the main appliance to create the queue managers as primary instances in the new disaster recovery configuration. You run the commands produced by that procedure on the recovery appliance to create secondary instances of the queue managers and synchronize them with the primary instance.

Procedure

1. Remove the existing recovery appliance from the network and power it down.
2. Power up the new recovery appliance and attach the supplied serial console cable to the console port.
3. Configure the ETH20 port. This is the port over which you will connect to the main appliance. See “Configuring Ethernet interfaces by using the command line” on page 105 for help with configuring the Ethernet ports.
4. Configure the MGT0 port with the same details that the appliance you are replacing was configured. (This connection enables you to send commands to the appliance from a remote workstation.)
5. Configure the remaining Ethernet ports as they were originally configured for the recovery appliance.
6. Connect the ETH20 port to the 10 Gb network that the ETH20 port of the recovery appliance is connected to.
7. Connect the MGT0 port to the 1 Gb network that the MGT0 port of the recovery port is connect to.
8. Ping the main appliance to check the connections.
9. Make other network connections as required to replicate the configuration of the original recovery appliance.
10. On the main appliance, complete the following steps:
    a. Enter MQ administration mode by typing `mqcli` on the command line.
    b. Enter the following command:

       ```
       crtdrprimary -m queue_manager -r new_appliance_name -i new_appliance_IP
       -p port
       ```

On successful completion, the command outputs a **crtdrsecondary** command.

c. Restart the queue manager:

```
strmqm queue_manager
```

11. On the recovery appliance, complete the following steps:
    a. Enter MQ administration mode by typing `mqcli` on the command line.
    b. Enter the **crtdrsecondary** command, exactly as output in step 10b. Synchronization of data from the main appliance to the recovery appliance begins.
    c. Check the progress of the synchronization by using the **status** command.

    ```
    status queue_manager
    ```
12. Repeat steps 10 and 11 for each queue manager.

Transfer from an existing high availability pair of appliances to a new pair of appliances

Follow this procedure to transfer queue managers from an existing high availability pair to a new high availability pair.

The following diagram gives an overview of the procedure.

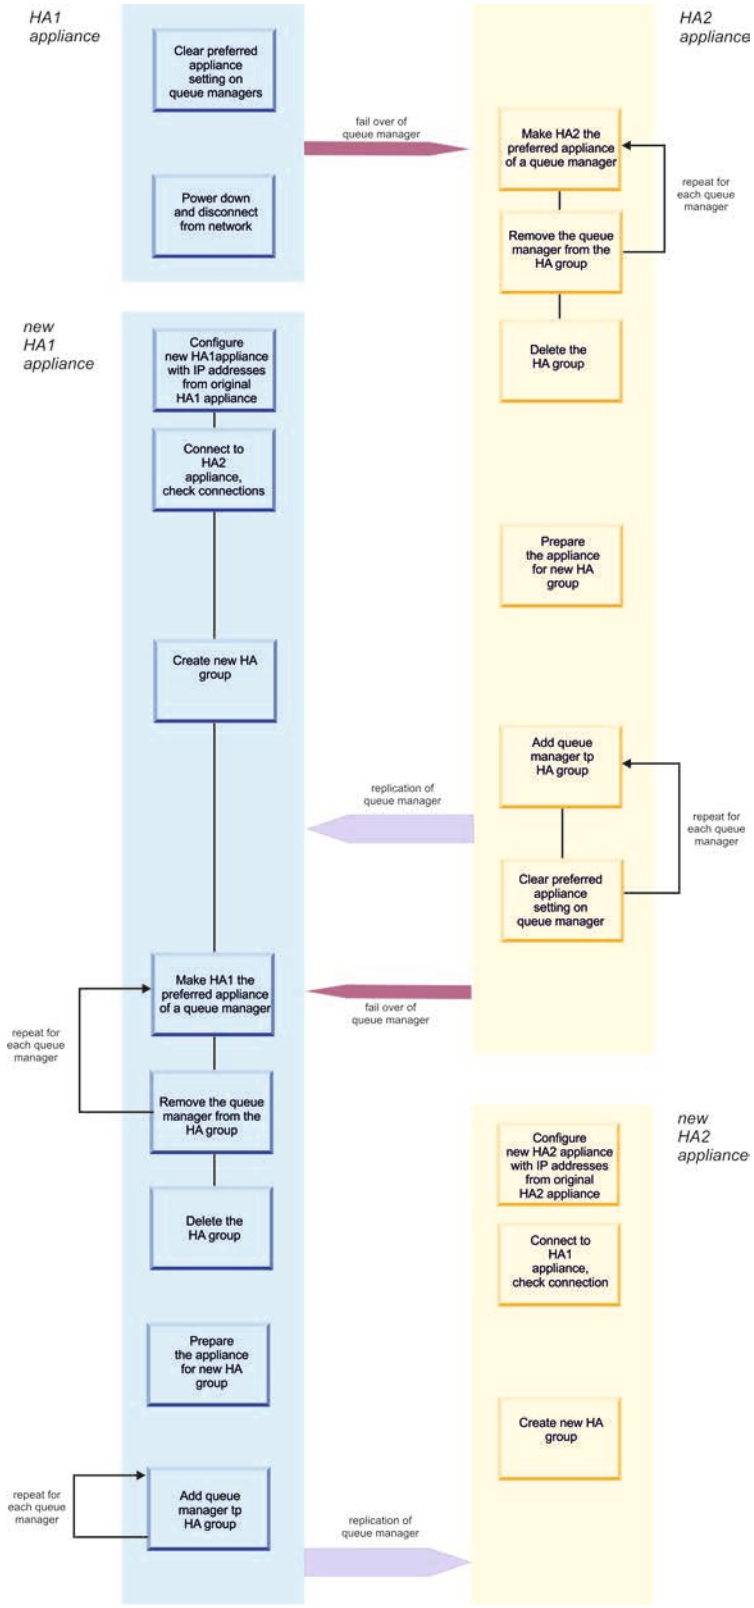

Figure 48. Transferring from an existing high availability group

Transferring from your first HA appliance

Follow this procedure to transfer queue managers from your first high availability appliance to the other appliance in the pair.

About this task

This procedure uses the following terminology:

- Appliance HA1 - the HA appliance named HA1. This is the first appliance that you upgrade. For the purposes of the description it is assumed that this is the preferred appliance for all your queue managers.
- Appliance HA2 - the second appliance in your HA pair. This is the second appliance that you upgrade.

The first part of the operation to transfer queue managers to a new, replacement appliance, is to run the queue managers on appliance HA2, and then deconstruct the HA group.

The appliance HA1 then has no active queue managers, you can disconnect the appliance and power it down.

Procedure

1. On appliance HA1, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Clear the preferred appliance setting from the queue manager that you want to work with:
      `clearhapreferred` *`queue_manager`*

      Repeat this command for each queue manager.
2. On appliance HA2, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Make the appliance the preferred appliance for the queue manager:
      `sethapreferred` *`queue_manager`*
   c. If the queue manager is running, stop it:
      `endmqm` *`queue_manager`*
   d. Remove the queue manager from the HA group:
      `sethagrp -e` *`queue_manager`*
   e. Repeat these steps for each queue manager.
   f. Delete the HA group:
      `dlthagrp`
3. Power down appliance HA1 and disconnect all the cables.

Preparing the new HA1 appliance

Follow these steps to prepare the new HA1 appliance.

About this task

To prepare the new appliance you power up the appliance and configure the Ethernet ports as they were configured for the original HA1 appliance.

Procedure

1. Power up the new HA1 appliance and attach the supplied serial console cable to the console port.

2. Configure the new appliance and set up ETH13, ETH17, and ETH21 to have the same configuration as the HA1 appliance that you are replacing. See "Configuring Ethernet interfaces by using the command line" on page 105 for help with configuring the Ethernet ports.
3. Configure the MGT0 port. Connect the port to a 1 Gb network, and configure it with the same details that the appliance you are replacing was configured. (This connection enables you to send commands to the appliance from a remote workstation.)
4. Configure the remaining Ethernet ports as they were originally configured for the main appliance.
5. Connect the new HA1 appliance to the existing HA2 appliance, as specified in the following table:

HA1 Appliance	HA2 Appliance
ETH21	ETH21
ETH17	ETH17
ETH13	ETH13

6. Make other network connections as required to replicate the configuration of the original HA1 appliance.
7. Check the 1 Gb and 10 Gb connections by ensuring that the HA1 appliance can ping the HA2 appliance and vice versa.

Transferring queue managers to the new HA1 appliance

Follow these steps to transfer a queue manager and associated data from the HA2 appliance to your new HA1 appliance.

About this task

After following the procedure to this point, your queue managers are currently located on appliance HA2, and are not part of an HA group. You now use HA commands to create a new HA group and replicate each queue manager to the new HA1 appliance.

Procedure

1. On the HA2 appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Enter the following command:
      ```
      prepareha -s secret_text -p IPaddress
      ```
      Where:
      - *secret_text* specifies a string that is used to generate a short-lived password. The password is used to set up the unique key for the two appliances.
      - *IPaddress* specifies the IP address that you have assigned to ETH13 on the new HA1 appliance.
2. On the new HA1 appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Enter the following command:
      ```
      crthagrp -s secret_text -p IPaddress
      ```
      Where:

- *secret_text* specifies the same string that you used with the **prepareha** command on the HA2 appliance.
- *IPaddress* specifies the IP address of ETH13 on the HA2 appliance.

3. On the HA2 appliance, add the queue managers back to the HA group. For each queue manager, ensure that it is stopped and enter the following command:

   ```
   sethagrp –i queue_manager
   ```

   Adding the queue manager to the HA group copies it to the HA1 appliance and restarts it running on the HA2 appliance.

Removing queue managers from your HA2 appliance

Follow this procedure to remove queue managers from the high availability configuration on appliance HA2 in preparation for replacing the HA2 appliance.

About this task

If you have followed the procedure to this point, you have a high availability pair comprising a new HA1 appliance and the original HA2 appliance. Queue managers are currently running on the HA2 appliance.

You now take steps to make the queue manager fail over to appliance HA1. Then you remove all the queue managers from the HA configuration, and delete the HA group.

The appliance HA2 now has no active queue managers, you can disconnect it and power it down.

Procedure

1. On appliance HA2, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Stop the first of the queue managers that you want to work with:
      ```
      endmqm queue_manager
      ```
   c. Clear the preferred appliance setting from the queue manager:
      ```
      clearhapreferred queue_manager
      ```

      Repeat steps b. and c. for each queue manager.
2. On appliance HA1, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Make the appliance the preferred appliance for the queue manager:
      ```
      sethapreferred queue_manager
      ```
   c. If the queue manager is running, stop it:
      ```
      endmqm queue_manager
      ```
   d. Remove the queue manager from the HA group:
      ```
      sethagrp -e queue_manager
      ```
   e. Repeat these steps for each queue manager.
   f. Delete the HA group:
      ```
      dlthagrp
      ```
3. Power down appliance HA2 and disconnect all the cables.

Preparing the new HA2 appliance

Follow these steps to prepare the new HA2 appliance.

About this task

To prepare the new appliance you power up the appliance and configure the Ethernet ports as they were configured for the original HA2 appliance.

Procedure

1. Power up the new HA2 appliance and attach the supplied serial console cable to the console port.
2. Configure the new appliance and set up ETH13, ETH17, and ETH21 to have the same configuration as the HA2 appliance that you are replacing. See "Configuring Ethernet interfaces by using the command line" on page 105 for help with configuring the Ethernet ports.
3. Configure the MGT0 port. Connect the port to a 1 Gb network, and configure it with the same details that the appliance you are replacing was configured. (This connection enables you to send commands to the appliance from a remote workstation.)
4. Configure the remaining Ethernet ports as they were originally configured for the main appliance.
5. Connect the new HA2 appliance to the HA1 appliance, as specified in the following table:

HA1 Appliance	HA2 Appliance
ETH21	ETH21
ETH17	ETH17
ETH13	ETH13

6. Make other network connections as required to replicate the configuration of the original HA2 appliance.
7. Check the 1 Gb and 10 Gb connections by ensuring that the HA2 appliance can ping the HA1 appliance and vice versa.

Creating a new HA group

Follow these steps to create a new HA group on the new HA1 and HA2 appliances.

About this task

After following the procedure to this point, your queue managers are currently located on appliance HA1, and are not part of an HA group. You now use HA commands to create a new HA group and replicate each queue manager to the new HA2 appliance.

Procedure

1. On the HA1 appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Enter the following command:

      ```
      prepareha -s secret_text -p IPaddress
      ```

      Where:

- *secret_text* specifies a string that is used to generate a short-lived password. The password is used to set up the unique key for the two appliances.
- *IPaddress* specifies the IP address that you have assigned to ETH13 on the new HA2 appliance.

2. On the new HA2 appliance, complete the following steps:
   a. Enter MQ administration mode by typing `mqcli` on the command line.
   b. Enter the following command:

      ```
      crthagrp -s secret_text -p IPaddress
      ```

      Where:
      - *secret_text* specifies the same string that you used with the **prepareha** command on the HA1 appliance.
      - *IPaddress* specifies the IP address of ETH13 on the HA1 appliance.
3. On the HA1 appliance, add the queue managers back to the HA group. For each queue manager, ensure that it is stopped and enter the following command:

   ```
   sethagrp –i queue_manager
   ```

   Adding the queue manager to the HA group copies it to the HA2 appliance and restarts it running on the HA1 appliance.

Chapter 8. Security

You must consider security requirements before and after you configure the IBM MQ Appliance.

User authentication and authorization

There are two types of user on the IBM MQ Appliance: appliance users, and messaging users. Appliance users are users that can administer the appliance and IBM MQ resources. Messaging users are users that can perform operations on messaging resources.

Appliance users

Appliance users can log in locally to the appliance and modify all appliance settings, including IBM MQ configuration and objects. Appliance users are created by using the appliance user commands or the **Appliance** section of the IBM MQ Appliance web UI.

Appliance users are stored in the internal user store. See “Configuring appliance users” on page 120 for guidance on working with appliance users and appliance user groups.

Messaging users

Messaging users can connect to queue managers remotely to send and receive messages. They can be authorized to remotely manage some aspects of queue managers by using client connections such as the IBM MQ Explorer. Messaging users are created by using user administration commands.

Messaging users can be stored in the internal user store, or in an external LDAP repository. (The internal user store is separate to the store used for appliance users.) The scalability of the internal store is limited, so in situations where many messaging users exist, an external LDAP repository provides better performance.

See “Administering messaging users” on page 190 for guidance on working with messaging users and messaging user groups.

See Overview of LDAP authorization in the IBM MQ documentation for guidance on using an external LDAP repository.

TLS certificate management

The IBM MQ Appliance supports the Transport Layer Security (TLS) protocol to provide link level security for message channels and MQI channels.

The IBM MQ Appliance supports the same levels of TLS as IBM MQ. However, on the IBM MQ Appliance you do not set up a key repository. When a queue manager is created on the appliance, a key repository is automatically created for that queue manager. The key repository is deleted when the queue manager is deleted. Each of the commands that are available for working with certificates require you to specify which queue manager the command is applied to, so that the correct key repository is used.

© Copyright IBM Corp. 2015, 2016

You can choose to use self-signed certificates, or CA certificates (issued by a trusted third party). Self-signed certificates can be used for test systems, but should not be used for production systems.

For self-signed certificates, you exchange copies of the public part of each certificate in order to establish the trust relationship between the end-points. The public part of the certificate is held in a file that you move between the end-points.

For more information about TLS, see Cryptographic security protocols in the IBM MQ documentation. For more information about TLS in IBM MQ, see SSL and TLS security protocols in the IBM MQ documentation.

Note: IBM MQ Appliance does not support ciphers other than FIPS ciphers. The SSLFIPS parameter for queue managers can still be enabled or disabled, but has no effect.

Working with self-signed certificates

A self-signed certificate can be used for testing a system while you are waiting for the officially signed certificate to be returned from the certificate authority (CA).

The self-signed certificate is generated by the queue manager. Self-signed certificates are not suitable for production use, for the following reasons:

- Self-signed certificates cannot be revoked, which might allow an attacker to spoof an identity after a private key is compromised. CAs can revoke a compromised certificate, which prevents its further use. CA-signed certificates are therefore safer to use in a production environment, though self-signed certificates are more convenient for a test system.
- Self-signed certificates never expire. This is both convenient and safe in a test environment, but in a production environment it leaves them open to eventual security breaches. The risk is compounded by the fact that self-signed certificates cannot be revoked.
- A self-signed certificate is used both as a personal certificate and as a root (or trust anchor) CA certificate. A user with a self-signed personal certificate might be able to use it to sign other personal certificates. In general, this is not true of personal certificates issued by a CA, and represents a significant exposure.

You can create a self-signed certificate by using the **createcert** command. The public part of the certificate data is extracted to a file. This public part of the certificate must be exchanged with any communicating partners, for example, IBM MQ clients or other queue managers. Exchanging the public part of the certificate establishes a trust relationship between the queue manager and the partner.

To add the public part of the certificate to other queue managers on an IBM MQ Appliance, you can use the **addcert** command.

For more information about certificates in IBM MQ, see Digital certificates.

Creating a self-signed certificate

You can create a self-signed certificate for a queue manager by using the **createcert** command on the command line.

About this task

A self-signed certificate can be used for testing a system while you are waiting for the officially signed certificate to be returned from the certificate authority (CA).

The self-signed certificate is generated by the queue manager. As it is signed by the queue manager, and it certifies the identity of the queue manager, it is not suitable for use in a production system.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. Create a self-signed certificate by entering the following command:
   ```
   createcert -m QMgrName -dn DistinguishedName -label LabelName
   ```
   Where:

   QMgrName
   Specifies the name of the queue manager that you want to create a certificate for.

   DistinguishedName
   Specifies the X.500 distinguished name that uniquely identifies the certificate.

   DistinguishedName is a string that is enclosed in double quotation marks. For example, `"CN=John Smith,O=IBM,OU=Test,C=GB"`. The CN, O, and C attributes are required.

   LabelName
   Specifies the name of the label that is associated with the certificate.

   If you do not specify a label name, a label is automatically generated. This label has the name `ibmwebspheremq<QMgrName>`, where *QMgrName* is the name of the queue manager in lowercase.

   Note: You can specify a number of optional parameters when you create a self-signed certificate. For more information, see "createcert" on page 376.
3. Optional: Exit the IBM MQ administration mode by entering the following command:
   ```
   exit
   ```

What to do next

After the self-signed certificate is generated, the public part of the certificate is extracted to a file that is placed in `mqpubcert://`.

Add the public part of the certificate to the key repository of any communicating partners. For example, you can add the public part of the certificate to IBM MQ clients or to other queue managers. To add the public part of the certificate to other queue managers on the IBM MQ Appliance, use the **addcert** command. For more information, see "Adding the public part of a self-signed certificate."

If you want to copy the certificate to another system to add it to the key repository of a communicating partner, you can use the **copy** command. For more information, see "Downloading certificates from the appliance" on page 276.

Adding the public part of a self-signed certificate

You can add the public part of a self-signed certificate to a queue manager by using the **addcert** command on the command line.

Before you begin

The certificate file that you want to add to the key repository must be on the appliance in the following location: `mqpubcert://`. You can upload a file to this location by using the **copy** command. For more information, see "Uploading certificates to the appliance" on page 277.

About this task

You must add the public part of the certificate to the key repositories of any partners that communicate with the queue manager for which the certificate was created. For example, the partners might be IBM MQ clients, or other queue managers. If the partner queue manager is running on the IBM MQ Appliance, use the **addcert** command to add the public part of the certificate to the key repository of the queue manager.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Add the public part of the self-signed certificate by entering the following command:

   ```
   addcert -m QMgrName -label Label -file         FileName
   ```

   Where:

   QMgrName
   Specifies the name of the queue manager that you want to add the public part of a certificate to.

   Label Specifies the label that is associated with the certificate.

   FileName
   Specifies the file that contains the public part of the certificate.

   The file must be available on the appliance. The file must be located in `mqpubcert://`

   Note: You can specify a number of optional parameters when you add the public part of a certificate. For more information, see "addcert" on page 375.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Working with CA-signed certificates

A certificate authority (CA) signed certificate is a certificate that is issued by a trusted third party. As it is signed by a certificate authority, the certificate provides assurance that the public key of the certificate truly belongs to the entity for which it is issued.

CA-signed certificates are suitable for use in production systems.

When you work with CA-signed certificates, there are two types of certificate that you use:

CA-signed certificate
The CA-signed certificate is the certificate that is digitally signed by the certificate authority. This certificate identifies the queue manager.

CA certificate
The CA certificate is the certificate that identifies the certificate authority that issued the CA-signed certificate.

If you do not have a CA-signed certificate, you can generate a certificate request by using the **createcertrequest** command. This command creates a new public-private key pair and a certificate request in the key repository. The certificate request is extracted to a file, which can then be sent to a certificate authority to be signed.

After you receive the signed certificate from the certificate authority, you must add the certificate to the key repository of the queue manager for which it was signed. You can add the CA-signed certificate to the key repository by using the **receivecert** command.

You must also add the CA certificate to the key repository of any communicating partners, for example, IBM MQ clients or other queue managers. To add the certificate to other queue managers on an IBM MQ Appliance, you can use the **addcert** command. If you are using a certificate chain, you must add each CA certificate in the chain to the key repository.

If the queue manager you are communicating with is not running on an appliance, you can add the certificate by following the instructions in the IBM MQ documentation. For more information, see Working with SSL or TLS.

When the CA-signed certificate expires, you can create a new request to send to the certificate authority by using the **recreatecertrequest** command.

For more information about certificate authorities and CA-signed certificates, see Digital certificates in the IBM MQ documentation.

Creating a certificate request

You can create a certificate request for a queue manager by using the **createcertrequest** command on the command line.

About this task

If you do not have a CA-signed certificate, you can generate a certificate request by using the **createcertrequest** command. This command creates a new public-private key pair and a certificate request in the key repository. The certificate request is extracted to a file, which can then be sent to a certificate authority to be signed.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. Create the certificate request by entering the following command:
   ```
   createcertrequest -m QMgrName -dn DistinguishedName -label LabelName
   ```
   Where:

   QMgrName
   Specifies the name of the queue manager that you want to create a certificate request for.

DistinguishedName
Specifies the X.500 distinguished name that uniquely identifies the certificate.

DistinguishedName is a string that is enclosed in double quotation marks. For example, `"CN=John Smith,O=IBM,OU=Test,C=GB"`. The CN, O, and C attributes are required.

LabelName
Specifies the name of the label that is associated with the certificate.

If you do not specify a label name, a label is automatically generated. This label has the name `ibmwebspheremq<QMgrName>`, where *QMgrName* is the name of the queue manager in lowercase.

Note: You can specify a number of optional parameters when you create a certificate request. For more information, see "createcertrequest" on page 378.

3. Optional: Exit the IBM MQ administration mode by entering the following command:
   `exit`

What to do next

The certificate request is extracted to a file that is placed into `mqpubcert:`.

You must send the certificate request file to the certificate authority to be signed. To copy the certificate to another system, you can use the **copy** command. For more information, see "Downloading certificates from the appliance" on page 276.

When you receive the signed certificate, you must add it to the key repository of the queue manager for which the request was created. See, "Receiving a CA-signed certificate."

Receiving a CA-signed certificate

You can receive a CA-signed certificate into the key repository of a queue manager by using the **receivecert** command on the command line.

Before you begin

The certificate file that you want to receive must be on the appliance in the following location: `mqpubcert://`. You can upload a file to this location by using the **copy** command. For more information, see "Uploading certificates to the appliance" on page 277.

About this task

After you receive the signed certificate from the certificate authority, you must add the certificate to the key repository of the queue manager for which it was signed. You can add the certificate to the key repository by using the **receivecert** command.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   `mqcli`
2. Receive the certificate by entering the following command:
   `receivecert -m QMgrName -file FileName`

Where:

QMgrName
Specifies the name of the queue manager for which you want to receive the certificate.

FileName
Specifies the name of the file that contains the certificate.

The file must be available on the appliance. The file must be located in `mqpubcert://`

Note: You can specify a number of optional parameters when you receive a certificate. For more information, see "receivecert" on page 387.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

What to do next

After the certificate is received, you must add the CA certificate that signed the CA-signed certificate to the key repository of any communicating partners. For example, you can add the public part of the certificate to IBM MQ clients or to other queue managers. You can add the CA certificate to other queue managers on the IBM MQ Appliance by using the **addcert** command. For more information, see "Adding a CA certificate" on page 272.

Listing certificate requests for a queue manager

You can list any outstanding certificate requests for a queue manager by using the **listcertrequest** command on the command line. An outstanding certificate request is a request where the CA-signed certificate has not been received. Ensure that each certificate request is sent to your certificate authority.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. List the certificate requests by entering the following command:

   ```
   listcertrequest -m QMgrName
   ```

   Where:

   QMgrName
   Specifies the name of the queue manager that you want to list the certificate requests for.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

What to do next

- You can view the details of each outstanding certificate request by using the **detailcertrequest** command. For more information, see "Viewing a certificate request for a queue manager" on page 272.
- After the certificate authority sends you the CA-signed certificate, you must receive the CA-signed certificate to the key repository. For more information, see "Receiving a CA-signed certificate" on page 270.

Viewing a certificate request for a queue manager

You can view the details of a specific outstanding certificate request for a queue manager by using the **detailcertrequest** command on the command line.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. View the details of the certificate request by entering the following command:

   ```
   detailcertrequest -m QMgrName -label        Label
   ```

   Where:

   QMgrName
   : Specifies the name of the queue manager for which the certificate request details are shown.

   Label
   : Specifies the label that is associated with the certificate request for which detailed information is shown.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Deleting a certificate request

You can delete an outstanding certificate request for a queue manager by using the **deletecertrequest** command on the command line.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Delete the certificate request by entering the following command:

   ```
   deletecertrequest -m QMgrName -label        Label
   ```

   Where:

   QMgrName
   : Specifies the name of the queue manager that you want to delete the certificate request from.

   Label
   : Specifies the label that is associated with the certificate request.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

What to do next

You can delete any associated certificate request files from `mqpubcert://` by using the **delete** command. For more information, see "Deleting certificates from the appliance" on page 278.

Adding a CA certificate

You can add a CA certificate to a queue manager by using the **addcert** command on the command line.

Before you begin

The certificate file that you want to add to the key repository must be on the appliance in the following location: `mqpubcert://`. You can upload a file to this location by using the **copy** command. For more information, see "Uploading certificates to the appliance" on page 277.

About this task

Any partners that communicate with the queue managers must have a copy of the CA certificate of the CA that signed the certificate of the queue manager. For example, the partners might be IBM MQ clients, or other queue managers. If the partner queue manager is running on the IBM MQ Appliance, use the **addcert** command to add the public part of the certificate to the key repository of the queue manager.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Add the CA certificate by entering the following command:

   ```
   addcert -m QMgrName -label Label -file         FileName
   ```

   Where:

   QMgrName
   Specifies the name of the queue manager that you want to add the certificate to.

   Label Specifies the label that is associated with the certificate.

   For more information about valid syntax for the certificate label, see http://www.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.sec.doc/q014340_.htm in the IBM MQ documentation.

   FileName
   Specifies the file that contains the certificate.

   The file must be available on the appliance. The file must be located in `mqpubcert://`

   Note: You can specify a number of optional parameters when you add the certificate. For more information, see "addcert" on page 375.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Renewing a CA-signed certificate

You can create a renewal request for a CA-signed certificate by using the **recreatecertrequest** command on the command line.

About this task

When a CA-signed certificate expires, you can create a renewal request to send to the certificate authority. The certificate request is extracted to a file, which can then be sent to a certificate authority.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. Create the certificate renewal request by entering the following command:
   ```
   recreatecertrequest -m QMgrName -dn DistinguishedName
   ```
   Where:

   QMgrName
   Specifies the name of the queue manager that you want to renew a certificate request for.

   DistinguishedName
   Specifies the X.500 distinguished name that uniquely identifies the certificate.

   DistinguishedName is a string that is enclosed in double quotation marks. For example, `"CN=John Smith,O=IBM,OU=Test,C=GB"`. The CN, O, and C attributes are required.

   Note: You can specify a number of optional parameters when you renew a certificate request. For more information, see "recreatecertrequest" on page 388.
3. Optional: Exit the IBM MQ administration mode by entering the following command:
   ```
   exit
   ```

What to do next

The certificate renewal request is extracted to a file that is placed into `mqpubcert://`.

You must send the file to the certificate authority to be renewed. To copy the file to another system, you can use the **copy** command. For more information, see "Downloading certificates from the appliance" on page 276.

When you receive the new CA-signed certificate, you must add it to the key repository of the queue manager for which the request was created. See, "Receiving a CA-signed certificate" on page 270.

Listing certificates for a queue manager

You can list the certificates that are stored in the key repository of a queue manager by using the **listcert** command on the command line. All certificates are listed, including self-signed certificates, CA-signed certificates, and CA certificates.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. List the certificates in one of the following ways:
   - List all the certificates by entering the following command:
     ```
     listcert -m QMgrName
     ```
     Where:

     QMgrName
     Specifies the name of the queue manager that you want to list the certificates for.

- List all the certificates and the valid-from and valid-to dates of those certificates by entering the following command:

  ```
  listcert -m QMgrName -expiry
  ```

  Where:

  QMgrName
  Specifies the name of the queue manager that you want to list the certificates for.

- List all the certificates and the valid-from and valid-to dates of certificates that expire within a specified number of days by entering the following command:

  ```
  listcert -m QMgrName -expiry Days
  ```

  Where:

  QMgrName
  Specifies the name of the queue manager that you want to list the certificates for.

  Days Specifies that the valid-from and valid-to dates are displayed for certificates that expire within that number of *days*.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Viewing a certificate for a queue manager

You can view the details of a certificate in the key repository of a queue manager by using the **detailcert** command on the command line.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. View the details of the certificate by entering the following command:

   ```
   detailcer"Viewing a certificate for a queue manager"t -m QMgrName -label Label
   ```

   Where:

   QMgrName
   Specifies the name of the queue manager for which the certificate details are shown.

   Label Specifies the label that is associated with the certificate for which detailed information is shown.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Deleting a certificate

You can delete a certificate from the key repository of a queue manager by using the **deletecert** command on the command line.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Delete the certificate by entering the following command:

   `deletecert -m QMgrName -label Label`

   Where:

   QMgrName
   Specifies the name of the queue manager that you want to delete the certificate from.

   Label Specifies the label that is associated with the certificate.
3. Optional: Exit the IBM MQ administration mode by entering the following command:

   `exit`

What to do next

You can delete any associated certificate files from `mqpubcert://` by using the **`delete`** command. For more information, see "Deleting certificates from the appliance" on page 278.

Managing certificates on the appliance

Certificate files and certificate request files for the IBM MQ Appliance are stored in `mqpubcert:`. You can upload, download, or delete files from this location. You can also back up and restore the key repository.

You can back up and restore the key repository by using the **`keybackup`** and **`keyrestore`** commands. For more information, see "Backing up a key repository" on page 202, and "Restoring a key repository" on page 203.

Downloading certificates from the appliance

Certificate files and certificate request files for the IBM MQ Appliance are stored in `mqpubcert:`. You can download files to a remote system from this location to add to communicating partners or to send to certificate authorities by using the **`copy`** command, or by using the IBM MQ Appliance web UI.

About this task

In this task, the files are downloaded by using the SCP protocol. However, you can download the files by using HTTP, HTTPS, SCP, or SFTP. For more information about using the **`copy`** command with these protocols, see "**`copy`**" on page 505.

Procedure

- To download a file by using the command line interface:
  1. Connect to the command line of the appliance as described in "Command line access" on page 97.
  2. Log in to the appliance as an administrator.
  3. Type `config` to enter configuration mode.
  4. Copy the file by typing the following command:

     `copy mqpubcert:certFileName scp://username@ipaddress:port//path/`

     Where:

     certFileName
     Specifies the name of the certificate file or certificate request file that you want to download.

username
Specifies the user name to log onto the remote system where the file is downloaded to.

ipaddress
The IP address and port of the remote system where the file is downloaded to.

path The file path on the remote system where the file is downloaded to.

- To download a file to your local system by using the IBM MQ Appliance web UI:
  1. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
  2. Open the mqpubcert folder.
  3. Click the certificate file name link to save the file to your local system (the exact method for saving the file depends on the type of browser that you use).

Example

The following example shows the download of the certificate for the queue manager `QM1`. The certificate has the default label `ibmwebspheremqqm1`. The file is downloaded to the `home` directory on the remote system `192.0.2.2:22`:

```
copy mqpubcert:ibmwebspheremqqm1 scp://user@192.0.2.2:22//home/
```

What to do next

If you want to delete the file at any time later on, you can delete the file from `mqpubcert://` by using the **`delete`** command. For more information, see "Deleting certificates from the appliance" on page 278.

Uploading certificates to the appliance

Certificate files and certificate request files for the IBM MQ Appliance are stored in `mqpubcert:`. You can upload files from a remote system to this location to add to the key repositories of the queue managers by using the **`copy`** command, or by using the IBM MQ Appliance web UI.

About this task

In this task, the files are uploaded by using the SCP protocol. However, you can upload the files by using HTTP, HTTPS, SCP, or SFTP. For more information about using the **`copy`** command with these protocols, see "**`copy`**" on page 505.

Note: For firmware versions before 8.0.0.5, if you copy a certificate from a UNIX system, you must change file permissions on the certificate before you copy it to the appliance. You must change the permissions to `-rw-r-----`. You can use the following command on the UNIX system to change the permissions:

```
chmod u=rw,g=r cert_file_name
```

Procedure

- To upload a certificate file by using the command line interface:
  1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
  2. Log in as a user in the administrators group.
  3. Type the following command to enter configuration mode:
     ```
     config
     ```

4. Copy the certificate file to the target appliance:
   `copy scp://username@ipaddress:port//path/ mqpubcert:///certFileName`
   Where:

   username
   Specifies the user name to log on to the remote system where the file is uploaded from.

   ipaddress
   The IP address and port of the remote system from where the file is uploaded

   path The file path on the remote system where the file is uploaded from.

   certFileName
   Specifies the name of the certificate file that you want to upload.
5. Type **exit** to leave config mode.

- To upload a certificate file from your local system by using the IBM MQ Appliance web UI:
  1. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
  2. Click **Actions** for the mqpubcert folder.
  3. Select **Upload files** from the **Actions** menu.
  4. Click **Browse**, and browse for the location of the certificate file on your local system.
  5. Click **Upload** to upload the certificate file to the mqpubcert directory on the appliance.

Example

The following example shows the upload of the certificate for the queue manager QM2. The certificate has the default label ibmwebspheremqqm2. The file is uploaded from the home directory on the remote system 192.0.2.2:22:

```
copy scp://user@192.0.2.2:22//home/ mqpubcert:///ibmwebspheremqqm2
```

What to do next

You can delete the file from mqpubcert: by using the **delete** command. For more information, see "Deleting certificates from the appliance."

Deleting certificates from the appliance

Certificate files and certificate request files for the IBM MQ Appliance are stored in mqpubcert:. You can delete files from this location by using the **delete** command.

Procedure

Delete a file by entering the following command:
`delete mqpubcert:certFileName`
Where:

certFileName
Specifies the name of the certificate file that you want to delete.

Example

The following example shows deleting the certificate ibmwebspheremq<qm1>:

```
delete mqpubcert:ibmwebspheremq<qm1>
```

Configuring certificates for IBM MQ Appliance web UI

You can configure the IBM MQ Appliance web UI to use certificates that you supply.

About this task

You use the appliance command line interface to configure the IBM MQ Appliance web UI to use your certificates.

To set up secure communication between a browser and the IBM MQ Appliance web UI and to handle certificates, you create an SSL server profile on the appliance. You import the required certificates and key file to the appliance, and create definition objects for them. The definition objects are used when you create an ID credentials (idcred) object for the appliance. The idcred is in turn used when you configure the SSL server profile. Finally, the SSL server profile is associated with your web management profile.

If you want to configure client validation, you import the certificates of the clients that are going to be allowed to connect. You then create definition objects for the certificates, which are used when you create a validation credential (valcred) object. The valcred object is in turn used when you configure the SSL server profile.

The example in this topic assumes that you have a signed certificate for the appliance. When making certificate requests for an appliance, the CN part of the distinguished name must be the URL that you type to reach the web UI. For example, `myappliance1.ourcompany.com`. If you want to set up the profile to validate connecting clients, you also require the relevant client certificates.

By default the web management service listens on all of the appliance ports (`local address` set to 0.0.0.0). You can, however, configure the service so that it listens on an IP address or host alias of a specific port (and so limit access to the web UI - see "Changing the IBM MQ Appliance web UI IP address and port" on page 99).

Procedure

- To upload certificates to your appliance:
  1. Ensure that you have the following items:
     - A private key to access the appliance certificate.
     - The appliance certificate.
     - Client certificates (optional).
  2. Connect to the IBM MQ Appliance as described in "Command line access" on page 97.
  3. Log in as a user in the administrators group.
  4. Type the following command to enter configuration mode:

     ```
     config
     ```

  5. Upload the key and certificates to the appliance by using the `copy` command, for example:

     ```
     copy scp://username@otherserver//home/username/myappliance1key.pem cert:
     copy scp://username@otherserver//home/username/myappliance1.cer cert:
     copy scp://username@otherserver//home/username/client1.cer cert:
     copy scp://username@otherserver//home/username/client2.cer cert:
     copy scp://username@otherserver//home/username/client3.cer cert:
     ```

     You can also copy the certificates to your appliance by using the IBM MQ Appliance web UI, see "Uploading certificates to the appliance" on page 277.

- To create definition objects for the appliance certificate and key:
  1. From configuration mode, type `crypto` to enter crypto configuration mode.
  2. Create a crypto key definition for the private key that is used for generating the appliance certificate:

     ```
     key key_alias cert:///keyfile
     ```

     For example:

     ```
     key WebUiKey01 cert:///myappliance1key.pem
     ```
  3. Create a crypto certificate definition for the appliance:

     ```
     certificate cert_alias cert:///certfile
     ```

     For example:

     ```
     certificate WebUiCert01 cert:///myappliance1.cer
     ```
  4. Create a crypto credential definition for the appliance:

     ```
     idcred credential_name key_alias cert_alias
     ```

     For example:

     ```
     idcred WebUiCred01 WebUiKey01 WebUiCert01
     ```
- To create a crypto valcred definition for validating clients (this is optional):
  1. From the crypto configuration mode, create a certificate definition object for each of the client certificates that you have imported:

     ```
     certificate cert_alias cert:///certfile
     ```

     For example:

     ```
     certificate WebUiClientCert01 cert:///client1.cer
     certificate WebUiClientCert02 cert:///client2.cer
     certificate WebUiClientCert03 cert:///client3.cer
     ```
  2. Create a crypto valcred definition, specifying the certificate definitions for the client certificates:

     ```
     valcred valcred_name
     certificate cert_alias
     ```

     Repeat the **certificate** command to specify the certificate definition for every client certificate that you have uploaded. For example:

     ```
     valcred WebUIvalcred01
     certificate WebUIClientCert01
     certificate WebUIClientCert02
     certificate WebUIClientCert03
     ```
- To create an SSL server profile for the appliance:
  1. From the crypto configuration mode, enter the following commands:

     ```
     ssl-server SSL_Svr_Profile_name
     admin-state enabled
     idcred IDCred_name
     protocols TLSv1d2
     ```

     If you are specifying client validation, also enter:

     ```
     valcred ValCred_name
     request-client-auth on
     require-client-auth on
     send-client-auth-ca-list on
     ```

     For example:

```
ssl-server myappliance1
admin-state enabled
idcred WebUiCred01
protocols TLSv1d2
valcred WebUIvalcred01
request-client-auth on
require-client-auth on
send-client-auth-ca-list on
```

- To save all the changes you have made in crypto configuration mode:
  1. Type `exit` to leave crypto configuration mode.
  2. Type `write mem` to save your configuration changes.
- To associate the SSL server profile with the web UI:
  1. From configuration mode, type `web-mgmt` to enter web management configuration mode.
  2. Enter the following command:

     `ssl-server` *`SSL_Svr_Profile_name`*

     For example:

     ```
     ssl-server myappliance1
     ```
- To save your web management configuration:
  1. Type `exit` to leave web-mgmt configuration mode.
  2. Type `write mem` to save your configuration changes.
  3. Type `exit` again to leave configuration mode.

Chapter 9. Monitoring and reporting

You can monitor the IBM MQ Appliance to understand how it is being used, and watch for potential problems. Monitoring provides a picture of the health of the IBM MQ Appliance.

Monitoring system resource usage

You can monitor the use of IBM MQ.

Monitoring system resource usage by using the status command

You can monitor the use of system resources on the appliance by using the **status** command on the command line.

About this task

You can use the **status** command to view the following information about the system resources on the appliance:

- The size and usage of the system memory
- The CPU usage of the system
- The size and usage of the internal disk
- The size and usage of the system volume
- The number of FDCs and the disk space used
- The disk space used by trace

You can use the **status** command to view the following information about the system resources that are used by a queue manager:

- The queue manager name
- The queue manager status
- An estimate of the CPU usage of the queue manager
- An estimate of the memory usage of the queue manager.
- The amount of the queue manager file system used by the queue manager

For a high availability queue manager, the following additional information can be viewed:

- The file system size for the queue manager
- The replication status of the queue manager
- The preferred appliance for the queue manager
- Whether a partitioned situation has been detected, and if it has, the amount of 'out-of-sync' data held.

For a disaster recovery queue manager, the following additional information can be viewed:

- The disaster recovery role (primary or secondary)
- The disaster recovery status
- The percentage complete if synchronization is in progress

© Copyright IBM Corp. 2015, 2016

- The estimated time to completion if synchronization is in progress
- The amount of out-of-sync data if the disaster recovery system is partitioned

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. View either the resource usage for the entire appliance, or for a specific queue manager:
   - To view the resource usage for the entire appliance, enter the following command:

     ```
     status
     ```

   - To view the resource usage for a specific queue manager on the appliance, enter the following command:

     `status` *qMgrName*Where:

     qMgrName
     Specifies the name of the queue manager that you want to view the resource statistics for.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the command to view the resource usage for a queue manager `QM1`:

```
status QM1
```

This command results in the following output:

```
QM(QM1)                    Status(Running)
CPU:                       29.89%
Memory:                    336MB
QMgr filesystem:           137MB used, 4.0GB allocated [3%]
```

Monitoring system resource usage by using the amqsrua command

You can use the **amqsrua** command to query metadata that is related to the system resource usage of a queue manager.

About this task

The **amqsrua** command reports metadata that is published by queue managers. This data can include information about the CPU, memory, and disk usage. You can also see data equivalent to the STATMQI PCF statistics data. The data is published every 10 seconds and is reported while the command runs.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Query the meta data by entering the following command:

   `amqsrua -n` *MaxPubs* `-m` *QMgrName*

Where:

MaxPubs
Specifies how many reports are returned before the command ends. The command publishes data every ten seconds, so if you enter a value of 50, the command returns 50 reports over 500 seconds.

If you do not specify this parameter, the command runs until either an error occurs, or the queue manager shuts down.

QMgrName
Specifies the name of the queue manager that you want to query. The queue manager must be running.

If you do not specify a queue manager name, the default queue manager is used.

3. From the list of options, enter the class of data you want returned. The class is case-sensitive. The following options are available:

CPU Returns information about CPU usage.

DISK Returns information about disk usage.

STATMQI
Returns information about MQI usage.

STATQ
Returns information about per-queue MQI usage.

4. From the list of options, enter the type of data you want returned. The type is case-sensitive. The following options are available:
   - For the **CPU** class:

     SystemSummary
     Returns information about CPU performance across the platform.

     QMgrSummary
     Returns information about CPU performance by the queue manager.
   - For the **DISK** class:

     SystemSummary
     Returns information about disk usage across the platform.

     QMgrSummary
     Returns information about disk usage by running queue managers.

     Log Returns information about disk usage by the queue manager recovery log.
   - For the **STATMQI** class:

     CONNDISC
     Returns information about calls to MQCONN and MQDISC.

     OPENCLOSE
     Returns information about calls to MQOPEN and MQCLOSE.

     INQSET
     Returns information about calls to MQINQ and MQSET.

     PUT Returns information about calls to MQPUT.

     GET Returns information about calls to MQGET.

     SYNCPOINT
     Returns information about syncpoint.

SUBSCRIBE
Returns information about calls to MQSUB.

PUBLISH
Returns information about subscribe requests.

- For the STATQ class:

  OPENCLOSE
  Returns information about calls to MQOPEN and MQCLOSE.

  INQSET
  Returns information about calls to MQINQ and MQSET.

  PUT Returns information about calls to MQPUT and MQPUT1.

  GET Returns information about calls to MQGET.

  After you have specified an option for the STATQ class, the appliance requests an object name. Specify the name of the queue that you want information for.

5. Optional: When **amqsrua** finishes, exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example shows the result of using **amqsrua** to view CPU performance data for the running queue manager over a 20-second period:

```
mqa(mqcli)# amqsrua -n 2 -m ASH
CPU : Platform central processing units
DISK : Platform persistent data stores
STATMQI : API usage statistics
STATQ : API per-queue usage statistics
Enter Class selection
==> CPU
SystemSummary : CPU performance - platform wide
QMgrSummary : CPU performance - running queue manager
Enter Type selection
==> QMgrSummary
Publication received PutDate:20151014 PutTime:09175398
User CPU time - percentage estimate for queue manager 0.02%
System CPU time - percentage estimate for queue manager 0.04%
RAM total bytes - estimate for queue manager 200MB

Publication received PutDate:20151014 PutTime:09180405
User CPU time - percentage estimate for queue manager 0.00%
System CPU time - percentage estimate for queue manager 0.00%
RAM total bytes - estimate for queue manager 200MB

mqa(mqcli)#
```

Monitoring system resource usage by using the IBM MQ Console

You use the Chart widget in the IBM MQ Console to view monitoring data for the IBM MQ Appliance.

About this task

You add a Chart widget to your dashboard and then configure it to monitor a particular aspect of appliance resource usage. You can create many instances of the Chart widget to display different data. The data is displayed in a chart format.

Data is collected at ten-second intervals. The X-axis of the chart displays a timeline. The Y-axis displays units appropriate to the resource that you are viewing. The Y-axis is dynamically resized to accommodate the data that is returned.

You add a Chart widget to your dashboard by clicking **Add Chart Widget**.

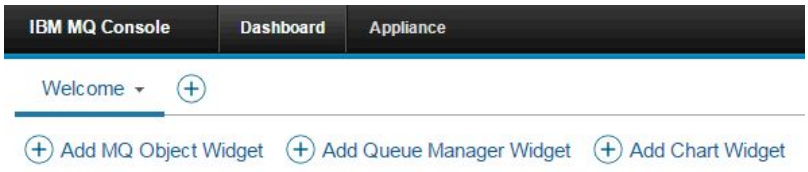

You configure the Chart widget by clicking the configure icon in the title bar . You must have at least one running queue manager before you can configure a chart widget.

After you configure the widget, there is a short delay before data is displayed in the chart. Data is displayed along a time axis as shown in the following example:

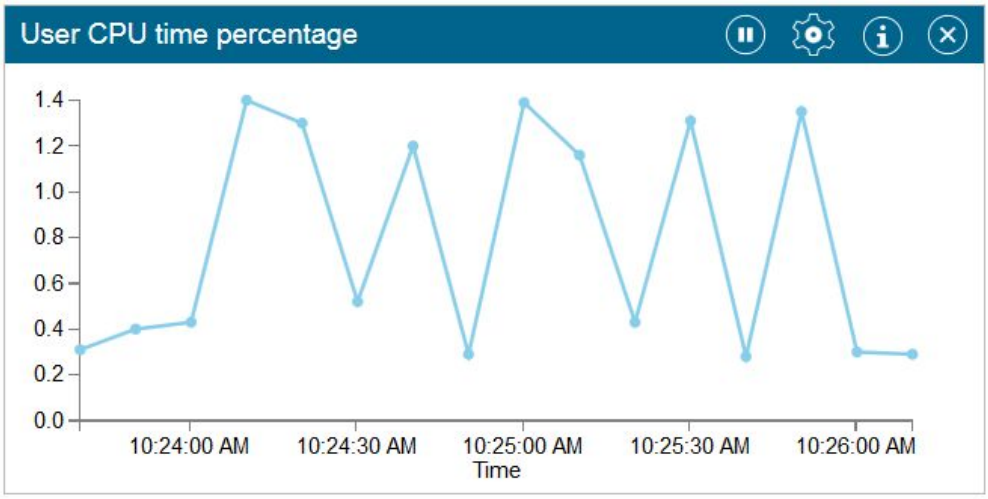

Each data point represents the end of the ten-second period over which the data is collected. You can hover over data points in the chart to see detailed information as shown in the following example:

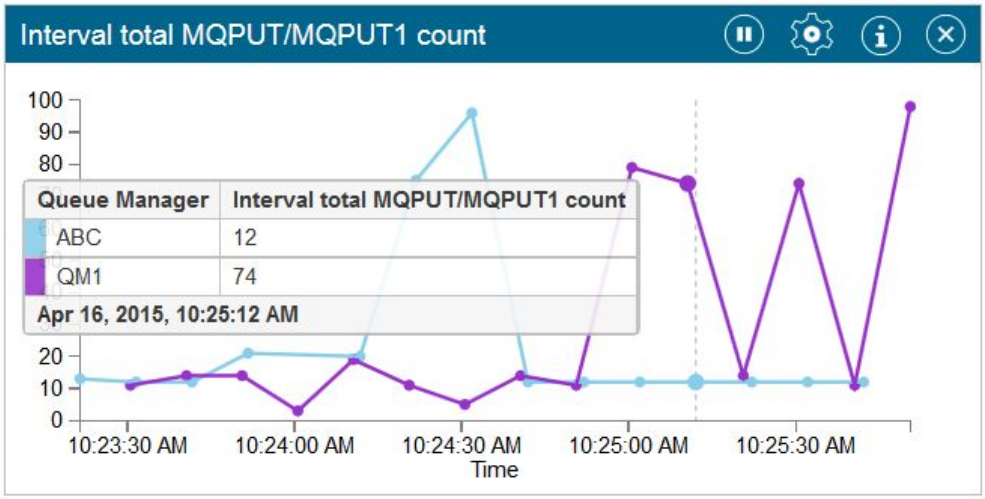

View the video for a demonstration of using the chart widget to monitor a queue manager running on an IBM MQ Appliance.

Using Charts with the IBM MQ Appliance and MQ Console

Procedure

1. Select the class of the resource that you want to monitor:

   Platform central processing units
   Monitor the usage of the CPUs on the IBM MQ Appliance.

   Platform persistent data stores
   Monitor the use of disk resource on the IBM MQ Appliance.

   API usage statistics
   Monitor API calls on the IBM MQ Appliance.

   API per-queue usage statistics
   Monitor API calls on the IBM MQ Appliance by individual queues. You specify the queue name when you choose this class.

2. Select the type of resource within the chosen class that you want to monitor:

Table 26. Resource types

Class	Type	Description
Platform central processing units	CPU performance – platform wide	Select this type to view performance data for the CPUs and memory in the appliance.
	CPU performance – running queue manager	Select this type to view performance data for the CPUs and memory that is related to the queue managers that you are monitoring. A queue manager must be running for you to monitor it. If you are monitoring results from more than one queue manager, different colors are used to distinguish the performance data in the chart.

Table 26. Resource types (continued)

Class	Type	Description
Platform persistent data stores	Disk usage – platform wide	Select this type to view performance data for global disk usage within the appliance.
	Disk usage - running queue managers	Select this type to view performance data for the disk usage that is related to the queue managers that you are monitoring. A queue manager must be running for you to monitor it. If you are monitoring results from more than one queue manager, different colors are used to distinguish the performance data in the chart.
	Disk usage - queue manager recovery log	Select this type to view data on how disk storage is currently being used for the recovery log of each queue manager that you are monitoring.
API usage statistics	MQCONN and MQDISC	Select this type to view data on MQCONN and MQDISC calls.
	MQOPEN and MQCLOSE	Select this type to view data on MQOPEN and MQCLOSE calls.
	MQINQ and MQSET	Select this type to view data on MQINQ and MQSET calls.
	MQPUT	Select this type to view data on MQPUT-related calls.
	MQGET	Select this type to view data on MQGET-related calls.
	commit and rollback	Select this type to view information about the use of syncpoints by the queue manager.
	subscribe	Select this type to view data related to MQSUB calls.
	publish	Select this type to view data about published messages.
API per-queue usage statistics	MQOPEN and MQCLOSE	Select this type to view data on MQOPEN and MQCLOSE calls for the specified queue.
	MQINQ and MQSET	Select this type to view data on MQINQ and MQSET calls for the specified queue.
	MQPUT and MQPUT1	Select this type to view data on MQPUT-related and MQPUT1-related calls for the specified queue.
	MQGET	Select this type to view data on MQGET-related calls for the specified queue.

3. Select the element for the type of resource monitoring that you selected:

Table 27. Elements for Appliance CPU types

Type	Element	Description
CPU performance – platform wide	User CPU time percentage	Shows the percentage of CPU busy in user state.
	System CPU time percentage	Shows the percentage of CPU busy in system state.
	CPU load – one minute average	Shows the UNIX load average over one minute.
	CPU load – five minute average	Shows the UNIX load average over five minutes.
	CPU load – fifteen minute average	Shows the UNIX load average over fifteen minutes.
	RAM free percentage	Shows the percentage of free RAM memory.
	RAM total bytes	Shows the total bytes of RAM configured.
CPU performance – running queue manager	User CPU time - percentage estimate for queue manager	Estimates the percentage of CPU use in user state for processes related to the queue managers that are being monitored.
	System CPU time - percentage estimate for queue manager	Estimates the percentage of CPU use in system state for processes related to the queue managers that are being monitored.
	RAM total bytes - estimate for queue managers	Estimates the total bytes of RAM in use by the queue managers that are being monitored.

Table 28. Elements for appliance data store resources

Type	Element	Description
Disk usage – platform wide	Appliance data - bytes in use	Shows the overall disk usage on the appliance.
	Appliance data - free space	Shows the overall free space on the appliance.
	System volume - bytes in use	
	System volume - free space	
	MQ trace file system - bytes in use	Shows the number of bytes of disk storage that are being used by the trace file system.
	MQ trace file system - free space	Shows the disk storage that is reserved for the trace file system that is free.
	MQ errors file system - bytes in use	Shows the number of bytes of disk storage that is being used by error data.
	MQ errors file system - free space	Shows the disk storage that is reserved for error data that is currently free.

Table 28. Elements for appliance data store resources (continued)

Type	Element	Description
	MQ FDC file count	Shows the current number of FDC files.
Disk usage - running queue managers	Queue Manager file system - bytes in use	Shows the number of bytes of disk storage that is used by queue manager files for the queue managers that you are monitoring.
	Queue Manager file system - free space	Shows the disk storage reserved for queue manager files that is free.
Disk usage - queue manager recovery log	Log - bytes in use	Shows the number of bytes of disk storage that is used for the recovery logs of the queue managers that you are monitoring.
	Log - bytes max	Shows the maximum bytes of disk storage that are configured to be used for queue manager recovery logs.
	Log file system - bytes in use	Shows the total number of disk bytes in use for the log file system.
	Log file system - bytes max	Shows the number of disk bytes that are configured for the log file system.
	Log - physical bytes written	Shows the number of bytes being written to the recovery logs.
	Log - logical bytes written	Shows the logical number of bytes written to the recovery logs.
	Log - write latency	Shows a measure of the latency when writing synchronously to the queue manager recovery log.

Table 29. Elements for MQI usage resources

Type	Element	Description
MQCONN and MQDISC	MQCONN/MQCONNX count	Shows the number of calls to MQCONN and MQCONNX.
	Failed MQCONN/ MQCONNX count	Shows the number of failed calls to MQCONN and MQCONNX.
	Concurrent connections - high water mark	Shows the maximum number of concurrent connections in the current statistics interval.
	MQDISC count	Shows the number of calls to MQDISC.
MQOPEN and MQCLOSE	MQOPEN count	Shows the number of calls to MQOPEN.
	Failed MQOPEN count	Shows the number of failed calls to MQOPEN.
	MQCLOSE count	Shows the number of calls to MQCLOSE.

Table 29. Elements for MQI usage resources (continued)

Type	Element	Description
	Failed MQCLOSE count	Shows the number of failed calls to MQCLOSE.
MQINQ and MQSET	MQINQ count	Shows the number of calls to MQINQ.
	Failed MQINQ count	Shows the number of failed calls to MQINQ.
	MQSET count	Shows the number of calls to MQSET.
	Failed MQSET count	Shows the number of failed calls to MQSET.
MQPUT	Interval total MQPUT/MQPUT1 count	Shows the number of calls to MQPUT and MQPUT1.
	Interval total MQPUT/MQPUT1 byte count	Shows the total bytes of data that is put by calls to MQPUT and MQPUT1.
	Non-persistent message MQPUT count	Shows the number of non-persistent messages that are put by MQPUT.
	Persistent message MQPUT count	Shows the number of persistent messages that are put by MQPUT.
	Failed MQPUT count	Shows the number of failed calls to MQPUT.
	Non-persistent message MQPUT1 count	Shows the number of non-persistent messages that are put by MQPUT1.
	Persistent message MQPUT1 count	Shows the number of persistent messages that are put by MQPUT1.
	Failed MQPUT1 count	Shows the number of failed calls to MQPUT1.
	Put non-persistent message - byte count	Shows the number of bytes put in non-persistent messages.
	Put persistent message - byte count	Shows the number of bytes put in persistent messages.
	MQSTAT count	Shows the number of calls to MQSTAT.
	Failed MQSTAT count	Shows the number of failed calls to MQSTAT.
MQGET	Interval total destructive get - count	Number of messages removed from queues by MQGET.
	Interval total destructive get - byte count	Bytes of data removed from queues by MQGET.
	Non-persistent message destructive get - count	Number of non-persistent messages removed from queues by MQGET.

Table 29. Elements for MQI usage resources (continued)

Type	Element	Description
	Persistent message destructive get - count	Number of persistent messages removed from queues by MQGET.
	Failed MQGET - count	Shows the number of failed calls to MQGET.
	Got non-persistent messages - byte count	Shows a count of bytes of non-persistent messages returned to MQGET.
	Got persistent messages - byte count	Shows a count of bytes of persistent messages returned to MQGET.
	Non-persistent message browse - count	Shows a count of non-persistent messages that have been browsed.
	Persistent message browse - count	Shows a count of persistent messages that have been browsed.
	Failed browse count	Shows a count of failed message browses.
	Non-persistent message browse - byte count	Shows the number of bytes of non-persistent messages that have been browsed.
	Persistent message browse - byte count	Shows the number of bytes of persistent messages that have been browsed.
	Expired message count	Shows a count of expired messages.
	Purged queue count	Shows a count of queues that have been purged.
	MQCB count	Shows the number of calls to MQCB.
	Failed MQCB count	Shows the number of failed calls to MQCB.
	MQCTL count	Shows the number of calls to MQCTL.
	Failed MQCTL count	Shows the number of failed calls to MQCTL.
Commit and rollback	Commit count	Shows the number of calls to MQCMIT.
	Failed commit count	Shows the number of failed calls to MQCMIT.
	Rollback count	Shows the number of calls to MQBACK.
Subscribe	Create durable subscription count	Shows the number of calls to MQSUB to create durable subscriptions.
	Alter durable subscription count	Shows the number of calls to MQSUB to alter durable subscriptions.

Table 29. Elements for MQI usage resources (continued)

Type	Element	Description
	Resume durable subscription count	Shows the number of calls to MQSUB to resume durable subscriptions.
	Create non-durable subscription count	Shows the number of calls to MQSUB to create non-durable subscriptions.
	Alter non-durable subscription count	Shows the number of calls to MQSUB to alter non-durable subscriptions.
	Resume non-durable subscription count	Shows the number of calls to MQSUB to resume non-durable subscriptions.
	Failed create/alter/resume subscription count	Shows the number of failed calls to MQSUBRQ to create, alter, or resume subscriptions.
	Delete durable subscription count	Shows the number of calls to MQSUB to delete durable subscriptions.
	Delete non-durable subscription count	Shows the number of calls to MQSUB to delete non-durable subscriptions.
	Subscription delete failure count	Shows the number of calls to MQSUB to delete subscriptions.
	MQSUBRQ count	Shows the number of calls to MQSUBRQ
	Failed MQSUBRQ count	Shows the number of failed calls to MQSUBRQ
	Durable subscriber - high water mark	Shows the maximum number of durable subscriptions in the current statistics interval.
	Durable subscriber - low water mark	Shows the minimum number of durable subscriptions in the current statistics interval.
	Non-durable subscriber - high water mark	Shows the maximum number of non-durable subscriptions in the current statistics interval.
	Non-durable subscriber - low water mark	Shows the minimum number of non-durable subscriptions in the current statistics interval.
Publish	Topic MQPUT/MQPUT1 interval total	The number of messages put to topics.
	Interval total topic bytes put	The number of message bytes put to topics.

Table 29. Elements for MQI usage resources (continued)

Type	Element	Description
	Published to subscribers - message count	Shows the number of messages published to subscribers.
	Published to subscribers - byte count	Shows the byte count of messages published to subscribers.
	Non-persistent - topic MQPUT/MQPUT1 count	Shows the number of non-persistent messages put to topics.
	Persistent - topic MQPUT/MQPUT1 count	Shows the number of persistent messages put to topics.
	Failed topic MQPUT/MQPUT1 count	Shows the number of failed attempts to put to a topic.

Table 30. Elements for MQI per-queue usage resources

Type	Element	Description
MQOPEN and MQCLOSE	MQOPEN count	Shows the number of calls to MQOPEN.
	MQCLOSE count	Shows the number of calls to MQCLOSE.
MQINQ and MQSET	MQINQ count	Shows the number of calls to MQINQ.
	MQSET count	Shows the number of calls to MQSET.
MQPUT and MQPUT1	MQPUT/MQPUT1 count	Shows the number of calls to MQPUT and MQPUT1.
	MQPUT byte count	Shows the total bytes of data that is put by calls to MQPUT and MQPUT1.
	MQPUT non-persistent message count	Shows the number of non-persistent messages that are put by MQPUT.
	MQPUT persistent message count	Shows the number of persistent messages that are put by MQPUT.
	MQPUT1 non-persistent message count	Shows the number of non-persistent messages that are put by MQPUT1.
	MQPUT1 persistent message count	Shows the number of persistent messages that are put by MQPUT1.
	non-persistent byte count	Shows the number of bytes put in non-persistent messages.
	persistent byte count	Shows the number of bytes put in persistent messages.
	queue avoided puts	

Table 30. Elements for MQI per-queue usage resources (continued)

Type	Element	Description
	queue avoided bytes	
	lock contention	
MQGET	MQGET count	
	MQGET byte count	
	destructive MQGET non-persistent message count	Number of non-persistent messages removed from the queue by MQGET.
	destructive MQGET persistent message count	Number of persistent messages removed from the queue by MQGET.
	destructive MQGET non-persistent byte count	Shows a count of bytes of non-persistent messages returned to MQGET.
	destructive MQGET persistent byte count	Shows a count of bytes of persistent messages returned to MQGET.
	MQGET browse non-persistent message count	Shows a count of non-persistent messages that have been browsed.
	MQGET browse persistent message count	Shows a count of persistent messages that have been browsed.
	MQGET browse non-persistent byte count	Shows the number of bytes of non-persistent messages that have been browsed.
	MQGET browse persistent byte count	Shows the number of bytes of persistent messages that have been browsed.
	messages expired	Shows a count of expired messages.
	queue purged count	Shows a count of queues that have been purged.
	average queue time	
	queue time	

4. Select a queue manager to monitor, and specify the color to display information in for that queue manager. Click **Add** to add more queue managers. You can specify up to five queue managers.
5. Optionally, specify a title for your Chart widget.
6. Click **Save**.

Monitoring the appliance by using the show command

You can monitor specific aspects of the operation of the appliance by using the **show** command on the command line.

About this task

You can use the show command to view information about how an aspect of the appliance is configured or to monitor aspects of the appliance operation. The

argument specifies which information you view. The show command is available at login, and in most configuration modes.

Procedure

- Enter the command:

  `show status_provider`

  Where *status_provider* identifies the status that the command displays.
- The available values for *status_provider* are listed in the following table:

Command	Description
"**show clock**" on page 577	Displays the current time and appliance uptime.
"**show cpu**" on page 578	Displays average CPU utilization data for the last 10 seconds, 1 minute, 10 minutes, 1 hour, and 1 day.
"**show ethernet-counters**" on page 578	Displays the statistics about Ethernet devices.
"**show ethernet-mau**" on page 578	
"**show ethernet-mii-registers**" on page 578	
"**show file**" on page 579	Displays a specified printable file.
"**show firmware-version**" on page 580	Displays the current firmware version, without image type and installation date.
"**show firmware**" on page 579	Displays the current firmware version, with image type and installation date.
"**show ipaddress**" on page 580	Provides IP address information about interfaces.
"**show link-aggregation-member-status**" on page 581	Lists members in link aggregation interfaces.
"**show link-aggregation-status**" on page 582	Provides statistics for aggregate interfaces.
"**show link**" on page 583	Provides status about all interfaces on the appliance.
"**show load**" on page 584	Displays task level system usage.
"**show log**" on page 584	Displays the appliance default log.
"**show logging**" on page 585	Displays a specified appliance log.
"**show loglevel**" on page 586	Displays the log-level for logging targets.
"**show ndcache**" on page 586	
"**show network-interface**" on page 587	Shows generic status of all network interfaces on the appliance.
"**show ntp-refresh**" on page 588	Lists the refresh status for the current NTP server.
"**show raid-array**" on page 589	Displays the status of the RAID array.
"**show raid-battery-module**" on page 589	Displays the information about the battery backup unit of the RAID controller.
"**show raid-logical-drive**" on page 590	Displays the status of the RAID logical drive.
"**show raid-physical-drive**" on page 591	Displays the status of the RAID physical drive.
"**show raid-ssd**" on page 592	Displays the estimated remaining life of the solid state disks
"**show route**" on page 592	Shows the routing table.

Command	Description
"**show sensors-current**" on page 593	Displays the values for sensors that read electrical current.
"**show sensors-fans**" on page 593	Displays the values for sensors that read the speed of the fans.
"**show sensors-other**" on page 594	Displays the status of sensors that have true or false values.
"**show sensors-temperature**" on page 595	Displays the values for sensors that read temperatures.
"**show sensors-voltage**" on page 596	Displays the values for sensors that read voltage.
"**show services**" on page 596	
"**show services-memory**" on page 596	
"**show system**" on page 597	Displays the System Settings configuration.
"**show tcp-connections**" on page 597	Lists the number of TCP connections in specific states.
"**show tcp-table**" on page 598	Lists the current TCP connections.
"**show tcp**" on page 598	Lists the current TCP connections followed by the number of connections in each state.
"**show throughput**" on page 598	Displays interface-specific traffic statistics.
"**show time**" on page 599	Displays the current time and appliance uptime.
"**show users**" on page 599	Lists all users who are currently logged in to the appliance.
"**show version**" on page 599	Displays the version of the firmware and libraries.

Developing your own resource monitoring program

You can develop your own program to monitor system resources.

Each queue manager publishes resource usage data to topics. This data is consumed by subscribers to those topics. When a queue manager starts, the queue manager publishes a set of messages on meta-topics. These messages describe which resource usage topics are supported by the queue manager, and the content of the messages published to those topics. Administrative tools can subscribe to the meta data to discover what resource usage information is available, and on what topics, and then subscribe to the advertised topics.

The topic tree for the meta data has the following structure:

`$SYS/MQ/INFO/QMGR/`*`QMGR-NAME`*`/Monitor/`*`class`*`[/`*`instance`*`]/`*`type`*`]`

For a list of possible classes and types, see "Monitoring system resource usage by using the amqsrua command" on page 284.

The source code for the **amqsrua** program is provided as an IBM MQ sample. You can use this program as a guide for creating your own monitoring program. You can retrieve the source for the sample from an IBM MQ client installation. The source file is named `amqsruaa.c` and is located in the samples directory:

- On Linux and UNIX platforms, *MQ_INSTALLATION_PATH*`/samp/`
- On Windows platforms, *MQ_INSTALLATION_PATH*`\tools\c\Samples\`

The **amqsrua** program subscribes to MQ resource usage topics and formats the resulting published PCF data. The program source provides a basic example of how to request and consume this type of administrative data. The **amqsrua** program completes the following tasks:

- Creates a non-durable subscription to the topics identified by the input parameters.
- Calls MQGET repeatedly to get messages from the topics, and writes to stdout.
- Writes a message for each MQI reason (other than MQRC_NONE).
- Stops if there is a MQI completion code of MQCC_FAILED, or when the requested number of resource usage publications have been consumed.

When your program is ready, you must run it on an IBM MQ client that connects to the queue manager you are monitoring. See "Setting up a queue manager to accept client connections" on page 192.

Interpreting CPU monitoring information

You can use the CPU monitoring information to see how well utilized your appliance is.

You can view the CPU monitoring information graphically by using the IBM MQ Console, or you can view a text report by using the **status** command. Using the IBM MQ Console gives you an ongoing view.

Two types of CPU statistics are reported: time percentage and CPU load. You can view CPU time percentage information for the appliance as a whole, or for a specified queue manager. CPU load is reported for the appliance as a whole.

Time percentage - appliance wide

The **status** command reports the total CPU percentage used. You can configure the chart widget in the IBM MQ Console to display either percentage of system CPU usage, or percentage of user CPU usage.

Time percentage - specified queue manager

The **status** command reports the total CPU percentage used by the specified running queue manager. You can configure the chart widget in the IBM MQ Console to display either percentage of system CPU usage, or percentage of user CPU usage for the running queue managers. The chart widget plots a line for each running queue manager on the appliance.

CPU load

The system load is a measure of the amount of computational work that a computer system performs. The load average represents the average system load over a period of time. It appears in the form of three numbers that represent the system load during the last one-, five-, and fifteen-minute periods. Each process using or waiting for CPU (the ready queue or run queue) increments the load number by 1. Each process that terminates decrements it by 1.

The IBM MQ Appliance M2000A/M2001A contains 20 cores, and a load average of 20 corresponds to the machine being 100% utilized. A load average of 10 correlates with the appliance being 50% loaded, and a load average of 40 corresponds to the machine having double the work requested that it can process. (The IBM MQ

Appliance M2000B/M2001B contains 6 cores.)

Example `status` output

In the following example, the appliance is 75% CPU utilized, of which the queue manager (PERF0) is using approximately 7%.

```
mqa(mqcli)# status
Memory:                                 16297MB used, 189.1GB total [8%]
CPU:                                    75%
CPU load:                               6.62, 6.19, 5.07
Internal disk:                          786432MB allocated, 2979.5GB total [26%]
System volume:                          5270MB used, 14.7GB allocated [35%]
MQ errors file system:                  175MB used, 2 FDCs, 15.8GB allocated [1%]
MQ trace file system:                   3034MB used, 31.5GB allocated [9%]


mqa(mqcli)# status PERF0
QM(PERF0)                               Status(Running)
CPU:                                    7.29%
Memory:                                 209MB
Queue manager file system:              1230MB used, 63.0GB allocated [2%]
```

Application activity trace

Application activity trace produces detailed information about the behavior of applications that are connected to a queue manager.

Application activity trace traces the behavior of an application and provides a detailed view of the parameters that are used by an application as it interacts with IBM MQ resources. It also shows the sequence of MQI calls issued by an application.

The IBM MQ Appliance supports the IBM MQ V8 methods of collecting application activity trace data, with the important exception that you cannot configure the collection of data by directly editing the `mqat.ini` configuration file. See Application activity trace in the IBM MQ documentation for details of collecting and reading application activity trace data. The methods write activity trace PCF messages to the system queue SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE.

In addition to writing trace data to the system queue, the IBM MQ Appliance introduces a new method of subscribing to activity trace data written to special IBM MQ system topics. This method is described in the following topics.

Note that the IBM MQ Appliance does not support the use of exits. If you have previously used exits to trace application activity, you must switch to using application activity trace.

Subscriptions to application activity trace

You can subscribe to an IBM MQ system topic to collect application activity trace information.

You subscribe to a special IBM MQ system topic string that represents the activity to trace. Subscribing automatically generates activity trace data messages and publishes them to the subscription destination queue. If you delete the subscription, the generation of activity trace data stops for that subscription.

A subscription can trace activity on one of the following resources:

- A specified application
- A specified IBM MQ channel
- An existing IBM MQ connection

You can create multiple subscriptions, with different, or the same topic strings. Where you create multiple subscriptions with the same system activity trace topic strings, each subscription receives a copy of the activity trace data, and this might have adverse performance implications.

Enabling any level of activity trace might have adverse performance effects. The more subscriptions, or the more resources subscribed to, the greater the potential performance overhead. To minimize the overhead of collecting activity trace, the data is written to messages and delivered to the subscriptions asynchronously from the application activity itself. Often, multiple operations are written to a single activity trace data message. The asynchronous operation can introduce a delay between the application operation and the receipt of the trace data that records the operation.

Creating subscriptions to application activity trace

You can create subscriptions to specific topics to collect application activity trace data on the IBM MQ Appliance.

When a subscription is created against specific system topic strings, appropriate activity trace PCF data messages are automatically published to that subscription. For detailed information on subscribing to topics, see Publish/subscribe messaging in the IBM MQ documentation.

The topic strings have the format:

`$SYS/MQ/INFO/QMGR/`*`qmgr_name`*`/ActivityTrace/`*`resource_type`*`/`*`resource_identifier`*

Where:

- *qmgr_name* specifies the queue manager that the traced application is connected to. *qmgr_name* is the name of the queue manager with all trailing blank characters removed and any forward slash (/) characters replaced by an ampersand (&) character.
- *resource_type* specifies the type of resource data is being collected for, and is one of the following strings:
  - `ApplName` to specify an application. The request subscribes to all IBM MQ connections that have an application name that matches the one specified by the *resource_identifier*.
  - `ChannelName` to specify an IBM MQ channel.
  - `ConnectionId` to specify an IBM MQ connection.
- *resource_identifier* identifies the actual resource. The format depends on the resource type:
  - For a resource type of `ApplName`, the *resource_identifier* is the trailing part (the value that follows the last / or \) of the application name as seen by the queue manager, with any trailing blank characters removed. The value matches the ApplName value from the API exit context structure (MQAXC). The ApplName of a connection is returned as the APPLTAG value when you use the MQSC command **`DISPLAY CONN`**.
  - For a resource type of ChannelName, the *resource_identifier* is the name of the channel to be traced. If the channel name identifies an SVRCONN channel, all application activity for connected clients is traced. If the channel name

identifies a queue manager to queue manager channel, the incoming and outgoing messages are traced. The *resource_identifier* is the channel name with all trailing blank characters removed and any '/' characters replaced by a '&' character.

– For a resource type of `ConnectionId`, the *resource_identifier* is the unique connection identifier that is assigned to each connection. The connection identifier in the topic string is the full 24-byte value written as a hexadecimal string. This value is the concatenation of the EXTCONN followed by the CONN values that are returned from the MQSC command **DISPLAY CONN**.

You can use wildcards in a *resource_identifier* to match multiple resource identities in a single subscription. The wildcard can either be in the default topic style ('#' or '+') or in the character style ('*' or '?'). When you use the topic style wildcard, it cannot be combined with part of a resource name, it can be used only to match all possible applications, channels, or connections. The use of any wildcards increases the level of trace data that is generated, which can affect performance.

To subscribe to these topic strings, you must have "subscribe" authorization. System topics do not inherit authorizations from the root of the queue manager topic tree. A user must be granted access to an administered topic object at or deeper than the $SYS/MQ point in the topic tree. You can subscribe if you have access to the SYSTEM.ADMIN.TOPIC, although this grants access to all $SYS/MQ topic strings, not just the activity trace. To control access more specifically, new administered topic objects can be defined for deeper points in the tree, either for all activity trace or, for example, for a specific application name or channel name.

Examples

The following example shows a topic string for an application that is named `amqsput` running on a Windows system:

```
$SYS/MQ/INFO/QMGR/QMGR1/ActivityTrace/ApplName/amqsputc.exe
```

The following example shows a topic string for a channel:

```
$SYS/MQ/INFO/QMGR/QMGR1/ActivityTrace/ChannelName/SYSTEM.DEF.SVRCONN
```

The following example shows a topic string for a connection:

```
$SYS/MQ/INFO/QMGR/QMGR1/ActivityTrace/ConnectionId/414D5143514D475231202020202020206B576B5420000701
```

The following example shows a topic string that creates a subscription to trace data for all channels on queue manager QMGR1:

```
$SYS/MQ/INFO/QMGR/QMGR1/ActivityTrace/ChannelName/#
```

The following example shows a topic string that creates a subscription to trace data for applications with names that start with "amqs" (note that to use the "*" wildcard, the subscription must be created using the character wildcard model):

```
$SYS/MQ/INFO/QMGR/QMGR1/ActivityTrace/ApplName/amqs*
```

Application activity trace: subscriptions compared with central collection

The IBM MQ Appliance supports two methods of collecting application activity trace data. There are points of overlap and differences between the two methods.

- Creating a subscription enables activity trace. You do not have to set queue manager or application attributes as for central collection of trace data. However,

any explicit blocking of activity trace by disabling trace at queue manager or application levels also blocks activity trace from being delivered to any matching subscriptions.

- You cannot edit the `mqat.ini` file directly to configure central activity trace collection on the IBM MQ Appliance.

Using `amqsact` to view trace messages

You can use the **amqsact** program with the IBM MQ Appliance to generate and view trace messages.

The **amqsact** program is an IBM MQ sample. To use this sample with the IBM MQ Appliance, you must use the client-connected executable file, **amqsactc**. The executable file is located in the samples directory:

- On Linux and UNIX platforms, *MQ_INSTALLATION_PATH*`/samp/bin64`
- On Windows platforms, *MQ_INSTALLATION_PATH*`\tools\c\Samples\Bin64`

You can use **amqsact** in two ways:

Display mode
: Format and display activity trace data messages that are being delivered to SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE.

Dynamic mode
: Create a subscription to a set of resources and display the generated activity trace by running **amqsact**.

Display mode

By default, **amqsact** in display mode processes messages on SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE. You can override this behavior by specifying a queue name or topic string. Activity trace must be enabled by using one of the methods that are described in Collecting application activity trace information in the IBM MQ documentation. You can control the trace period that is displayed and specify whether the activity trace messages are removed or retained after display. In display mode, **amqsact** takes the following arguments:

-m *queue_manager_name*
: Required. Specify the queue manager that trace messages are collected for.

-q *queue_name*
: Display only trace messages that are related to the named queue.

-t *topic_string*
: Display only trace messages that are related to the named topic.

-b
: Specify that trace messages are retained after display.

-v
: Display trace messages in verbose mode.

-d *depth*
: The number of messages to display.

-w *timeout*
: Specify a timeout. If no trace messages appear in that period, **amqsact** exits.

-s *start_time*
: Use this argument with the -e argument to specify a time period. Trace messages from the specified time period are displayed.

-e *end_time*
Use this argument with the -s argument to specify a time period. Trace messages from the specified time period are displayed.

For example, the following command displays activity trace messages that are held on SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE, and deletes the messages after display:

```
amqsact –m QMGR1
```

The following command displays activity trace messages on the specified queue, SUB.QUEUE, and deletes the messages after display. Messages continue to be displayed until a period of 30 seconds with no new messages elapses. This command can, for example, be used with a subscription to an activity trace system topic string.

```
amqact –m QMGR1 –q SUB.QUEUE.1 –w 30
```

The following command displays in verbose format any activity trace data that is currently held on the SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE that occurred in the 20-minute period specified. Messages will remain on the queue after display.

```
amqsact –m QMGR1 –b -v –s 2014-12-31 23.50.00 –e 2015-01-01 00.10.00
```

Dynamic mode

You enable dynamic mode by specifying an application name, a channel name, or a connection identifier as an argument to **amqsact**. You can use wildcard characters in the name. In dynamic mode, activity trace data is enabled at the start of the sample by use of a non-durable subscription to a system topic. Collecting activity trace data stops when **amqsact** stops. You must specify a timeout for **amqsact** in dynamic mode. You can run multiple copies of **amqsact** concurrently, and each instance receives a copy of any activity trace data. In dynamic mode, **amqsact** takes the following arguments:

-m *queue_manager_name*
Required. Specify the queue manager that trace messages are collected for.

-w *timeout*
Required. Specify a timeout. If no trace messages appear in that period, **amqsact** exits.

-a *application_name*
Specify an application to collect messages for.

-c *channel_name*
Specify a channel to collect messages for.

-i *connection_id*
Specify a connection to collect messages for.

-v Display trace messages in verbose mode.

For example, the following command generates and displays activity trace messages for any connections that are made by applications that are named “amqsget.exe”. After 30 seconds of inactivity, the **amqsact** program ends, and no new activity trace data is generated.

```
amqsactc -m QMGR1 -w 30 -a amqsget.exe
```

The following command generates and displays activity trace messages for any connections that are made by applications that start with the text “amqs”. After 30 seconds of inactivity, the **amqsact** program ends, and no new activity trace data is generated.

```
amqsactc -m QMGR1 -w 30 -a amqs*
```

The following command generates and displays activity trace messages for any activity on the QMGR1.TO.QMGR2 channel. After 10 seconds of inactivity, the **amqsact** program ends, and no new activity trace data is generated.

```
amqsactc -m QMGR1 -w 10 -c QMGR1.TO.QMGR2
```

The following command generates and displays activity trace messages for any activity on any channels. After 10 seconds of inactivity, the **amqsact** program ends, and no new activity trace data is generated.

```
amqsactc -m QMGR1 -w 10 -c #
```

The following command generates and displays verbose activity trace messages for any activity on the existing IBM MQ connection that has a CONN of "6B576B5420000701", and an EXTCONN of "414D5143514D47523120202020202020". After a minute of inactivity, the **amqsact** program ends, and no new activity trace data is generated.

```
amqsactc -m QMGR1 -w 60 -i 414D5143514D475231202020202020206B576B5420000701 -v
```

Configuring trace levels

You configure trace levels for a queue manager on the IBM MQ Appliance by using the **setmqini** command.

You use the **setmqini** command to set values in the `mqat.ini` file for the queue manager. See “Adding a value to the configuration file” on page 245 for details of how to use the **setmqini** command.

You can set the following values for the AllActivityTrace stanza:

ActivityInterval
: Time interval in seconds between trace messages. Activity trace does not use a timer thread, so the trace message is not written at the exact instant that the time elapses, it is written when the first MQI operation is executed after the time interval elapses. If this value is 0, the trace message is written when the connection disconnects (or when the activity count is reached). Defaults to 1.

ActivityCount
: Number of MQI operations between trace messages. If this value is 0, the trace message is written when the connection disconnects (or when the activity interval elapses). Defaults to 100.

TraceLevel
: Amount of parameter detail that is traced for each operation. The description of individual operations details which parameters are included for each trace level. Set to LOW, MEDIUM, or HIGH. Defaults to MEDIUM.

TraceMessageData
: Amount of message data that is traced in bytes for MQGET, MQPUT, MQPUT1, and Callback operations. Defaults to 0.

StopOnGetTraceMsg
Can be set to ON or OFF. Defaults to ON.

SubscriptionDelivery
Can be set to BATCHED or IMMEDIATE. Determines whether the ActivityInterval and ActivityCount parameters are to be used when one or more activity trace subscriptions are present. Setting this parameter to IMMEDIATE results in the ActivityInterval and ActivtyCount values being overridden with effective values of 1 when the trace data has a matching subscription. Each activity trace record is not batched with other records from the same connection and instead delivered to the subscription immediately with no delay. The IMMEDIATE setting increases the performance overhead of collecting activity trace data. The default setting is BATCHED.

System topics for monitoring and activity trace

System topics in queue manager topic trees are used for resource monitoring and for application activity trace.

Each queue manager's topic tree contains the $SYS/MQ branch. The queue manager publishes to topic strings in this branch. An authorized user can subscribe to these topic strings to receive information on the queue manager and the activity on it. These system topics are used for both monitoring resources on the IBM MQ Appliance and for application activity trace. For more information on topic trees, see Topic Trees in the IBM MQ documentation.

The root of the $SYS/MQ branch is represented by the SYSTEM.ADMIN.TOPIC topic object. The $SYS/MQ branch of the topic tree is isolated from the rest of the topic tree in the following ways:

- A subscription that is made with wildcard characters at a point higher in the tree than $SYS/MQ does not match any topic string within the $SYS/MQ branch. The wildcard operation for SYSTEM.ADMIN.TOPIC is set to "Block" and cannot be modified. This limitation also applies when you use wildcard characters with the **runmqsc** command DISPLAY TPSTATUS to display nodes in the topic tree. To view topic nodes within the $SYS/MQ branch, start the topic string with $SYS/MQ. For example, use `$SYS/MQ/#` to see all nodes.
- A user must be authorized at or deeper than $SYS/MQ to be granted authority to use the $SYS/MQ topic tree. Authorization to subscribe to a topic string is based on authorization being granted for an administered topic object at or higher than the topic string in the topic tree. Authorizations that are granted at the very root (SYSTEM.BASE.TOPIC) would grant a user authority to all topic strings. However, in the case of the $SYS/MQ branch, access granted higher than $SYS/MQ does not apply to the $SYS/MQ topic strings.
- The $SYS/MQ branch of the topic tree is isolated from topic attributes set higher in the tree. The SYSTEM.ADMIN.TOPIC does not inherit any attributes from a topic object defined higher in the topic tree. For example, changing attributes of SYSTEM.BASE.TOPIC does not affect the behavior of the $SYS/MQ branch.

All topic strings that start with $SYS/MQ are reserved for use by IBM MQ. These topic strings have the following restrictions:

- You cannot enable multicast from the $SYS/MQ branch of the topic tree.
- Clustering is not supported for the $SYS/MQ branch.
- The proxy subscription mechanism cannot be set to "force".
- Applications cannot publish to a $SYS/MQ topic string.

- Publication and subscription scope defaults to the local queue manager only.
- The use of wildcard characters within subscription topic strings is restricted. No wildcard characters can be used at the following points:
  - $SYS/MQ/
  - $SYS/MQ/INFO
  - $SYS/MQ/INFO/QMGR
  - $SYS/MQ/INFO/QMGR/*queue_manager_name*
  - $SYS/MQ/INFO/QMGR/*queue_manager_name*/ActivityTrace

  Attempts to use wildcard characters at these points causes a subscription failure with the reason MQRC_ADMIN_TOPIC_STRING_ERROR.

Chapter 10. Troubleshooting

You can use the troubleshooting information to help you to diagnose and resolve problems that you experience with your IBM MQ Appliance.

There are a number of diagnostic tools that you can use to help you resolve problems:

- You can list or view the system error logs, queue manager error logs, and first failure data captures (FDCs) by using the **dspmqerr** command.
- You can start and stop trace, and you can download the generated trace files by using the **strmqtrc** and **endmqtrc** commands.
- You can start and stop trace, and you can download the generated trace files by using the IBM MQ Console.
- You can view information about return codes by using the **mqrc** command.
- You can gather diagnostic information to send to IBM support by using the **runmqras** command.
- You can configure, generate, and put a trace-route message into a queue manager network by using the **dspmqrte** command. For more information, see dspmqrte in the IBM MQ documentation.

Error logs

There are several types of error logs generated by the IBM MQ Appliance, including system error logs, queue manager error logs, and first failure data captures (FDCs).

System error logs

The system error logs contain information about errors that occur where a queue manager name is not known. For example, if there are problems in a listener or a TLS handshake, the information is logged in the system error log. These files have a file type of `LOG`.

Queue manager error logs

The queue manager error logs contain information about errors that occur on a particular queue manager. This information includes messages that are related to channels that belong to the queue manager, unless the queue manager is unavailable, or the queue manager name is unknown. In this case, channel related messages are recorded in the system error log. These files have a file type of `LOG`.

You can configure the queue manager error logs:

- You can restrict the maximum size of the log file.
- You can exclude particular messages from the log.
- You can prevent repeats of particular messages within a set time interval.

For more information about how to configure the queue manager error logs, see Queue manager error logs in the IBM MQ documentation.

© Copyright IBM Corp. 2015, 2016

First failure data captures (FDCs)

First failure data captures provide an automated snapshot of the system environment when an unexpected internal error occurs. This snapshot is used by IBM support personnel to provide a better understanding of the state of the system when the problem occurred. These files have a file type of `FDC`.

For more information about FDCs, see First Failure Support Technology (FFST) in the IBM MQ documentation.

Viewing the logs

You use the **dspmqerr** command to view all types of error logs. See the child topics for details of how to use the command for each log type. The command is based on the UNIX `less` command. You can take the following actions:

- Use the arrows keys to scroll up and down the logs.
- Use the page, space, or return keys for simple scrolling.
- Enter `q` to exit at any time
- Enter `h` to display full help while you view a log. The help lists further commands, for example, for searching for strings or jumping a set number of lines.

Note: Some controls (for example, those controls that manipulate file names) are disabled for security reasons. If you try to use these controls, you get the message `Command not available`.

Viewing system error log files

You can view the system error log files by using the **dspmqerr** command on the command line. The log files are displayed on the command line. You can choose to display the most recent error log file, or a specific log file.

About this task

The command is based on the UNIX `less` command. The `less` command provides controls for navigating the contents of a file, and you can use these controls when you view system error logs.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   ```
   mqcli
   ```
2. Choose which system error log file to view:
   - To display the most recent system error log file, enter the following command:
     ```
     dspmqerr -s
     ```
   - To list system error log files, enter the following command:
     ```
     dspmqerr -s -l
     ```
   - To display a specific system error log file, enter the following command:
     ```
     dspmqerr -s Filename
     ```
     Where:

     Filename
     Specifies the name of the system error log file to display.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following example displays the system error log file AMQERR02.LOG:

```
dspmqerr -s AMQERR02.LOG
```

Viewing queue manager error log files

You can view the log files generated by a queue manager either from the appliance command line, or from the IBM MQ Console.

Viewing queue manager error log files by using the command line

You can view the queue manager error log files by using the **dspmqerr** command on the command line. The log files are displayed on the command line. You can choose to display the most recent error log file, or a specific log file.

About this task

The command is based on the UNIX `less` command. The `less` command provides controls for navigating the contents of a file, and you can use these controls when you view system error logs.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Choose which queue manager error log file to view:
   - To display the most recent log file for a queue manager, enter the following command:

     ```
     dspmqerr -m QMgrName
     ```

     Where:

     QMgrName
     Specifies the name of the queue manager that the log file is associated with.
   - To list error log files for a queue manager, enter the following command:

     ```
     dspmqerr -l -m QMgrName
     ```

     Where:

     QMgrName
     Specifies the name of the queue manager that the log files are associated with.
   - To display a specific log file for a queue manager, enter the following command:

     ```
     dspmqerr -m QMgrName Filename
     ```

     Where:

     QMgrName
     Specifies the name of the queue manager that the log file is associated with.

Filename
Specifies the name of the system error log file to display.

3. Optional: Exit the IBM MQ administration mode by entering the following command:
   `exit`

Example

The following command lists all the error logs for the queue manager `QM1`:

```
dspmqerr -l -m QM1
```

Viewing queue manager error log files by using the console

You can view the queue manager error log files by using the IBM MQ Console. You can choose to display the most recent error log file, or a specific log file.

About this task

You can download the log file to your local computer, or you can view it.

Procedure

1. Start the IBM MQ Appliance web UI and click **MQ Console**.
2. In the Queue Manager widget, select the queue manager that you want to view logs for.
3. Select **More** > **View Logs** from the Queue Manager widget toolbar.
4. Select the log that you want to view or download, and click the download icon, .
5. You can view the file, or save it to your local disk, depending on the options offered by your web browser.

Viewing the first failure data captures

You can view the first failure data captures (FDCs) by using the **dspmqerr** command on the command line. The FDCs are displayed on the command line.

About this task

The command is based on the UNIX `less` command. The `less` command provides controls for navigating the contents of a file, and you can use these controls when you view system error logs.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:
   `mqcli`
2. Choose which FDC to view:
   - To list all the FDCs that are available to view on the appliance, enter the following command:
     `dspmqerr -f -l`
   - To display a specific FDC, enter the following command:
     `dspmqerr -f Filename`
     Where:

Filename
Specifies the name of the FDC to display.

3. Optional: Exit the IBM MQ administration mode by entering the following command:

   ```
   exit
   ```

Example

The following command displays the FDC AMQ12345.FDC:

```
dspmqerr -f AMQ12345.FDC
```

Deleting log files

You can periodically purge log files to prevent them taking too much disk space.

About this task

You use the **dltmqras** command to delete log files. You can delete all log files, or specify the type of log files to delete.

For each file deleted, a message in the form `File deleted: filename` is written to `MQSystem.log`. You can view `MQSystem.log` by using the **dspmqerr** command without parameters.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

   ```
   mqcli
   ```

2. Specify what types of file you want to delete:
   - `dltmqras -a` to delete all log files apart from queue manager log files
   - `dltmqras -d` to delete general diagnostics files
   - `dltmqras -e` to delete older error logs. The current error log (`MQSystem.log`) is not deleted.
   - `dltmqras -f` to delete FDC files
   - `dltmqras -h` to delete HA log files
   - `dltmqras -m qmname` to delete service tool output for the specified queue managers
   - `dltmqras -p` to delete files in the `mqtemporary:` location
   - `dltmqras -t` to delete trace files
   - `dltmqras -w` to delete console log files
3. Optionally specify `-y` so that you are not prompted to confirm deletion. For example:

   ```
   dltmqras -a -y
   ```

Downloading error logs

You can view, delete, and download log files from the IBM MQ Appliance.

Directory structures on the appliance are accessible in the form of URIs. There is a dedicated URI, `mqerr`, for accessing IBM MQ logs. Use this URI to access queue manager logs, FDC files, and the system log.

You enter the commands to download files on the IBM MQ Appliance command line. Connect to the appliance as described in "Command line access" on page 97. Log in as an administrative user and type the command:

```
config
```

You can also view the URIs and view, delete, and download log files by using the IBM MQ Appliance web UI.

Listing the log directory

To list the contents of the log directory by using the command line, enter the following command:

```
dir mqerr:
```

To list the contents of the log directory by using the IBM MQ Appliance web UI:

1. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
2. Open the mqerr folder.
3. Select the log file that you want to view.

Deleting a log file

To delete a log file by using the command line, enter the following command:

```
delete mqerr:///logfile
```

For example:

```
delete mqerr:///MQSystem.log
```

To delete a log file by using the IBM MQ Appliance web UI:

1. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
2. Open the mqerr folder.
3. Select the log file that you want to delete.
4. Click **Delete**.

Downloading a log file

To download a log file from the appliance to your local system by using the command line, enter the following command:

```
copy mqerr:///logfile scp://username@ipaddress/[/]directory/
```

For example:

```
copy mqerr:///MQSystem.log scp://me@mycomputer//logfiles/
```

To download a log file by using the IBM MQ Appliance web UI.

1. Start the IBM MQ Appliance web UI, and click the **File Management** tab.
2. Open the mqerr folder.
3. Click the log file name link to save the file to your local system (the exact method for saving the file depends on the type of browser that you use).

Reason codes

IBM MQ Appliance has some new error reason codes in addition to the error codes listed in the IBM MQ documentation.

MQRCCF_TOPIC_RESTRICTED
This error can occur when you create or modify a topic object. One or more attributes of the topic object are not supported on an IBM MQ administrative topic. Modify the configuration to adhere to the documented restrictions.

MQRC_ADMIN_TOPIC_STRING_ERROR
This error can occur when calling MQSUB or MQOPEN. Publishing to an IBM MQ admin topic string that starts with $SYS/MQ/ is not permitted. When you subscribe to an IBM MQ admin topic string, the use of wildcards is restricted, see "System topics for monitoring and activity trace" on page 306 for details.

Event logs

Log targets capture messages that are posted by the various objects and services that are running on the appliance. These are appliance-specific objects and services, IBM MQ logging is separate.

An appliance supports a maximum of 500 log targets.

Log targets capture events that occur because of some internal process or hardware status change.

Different types of log targets might include one or more of the following capabilities:

- Archive files through rotation or upload
- Forward messages to remote servers

Types of log target

Target types enable additional capabilities that include rotating files and sending files to remote servers.

The following types of log targets are available.

Cache Writes log entries to memory.

Console
Writes log entries to the screen with Telnet, SSH, or command-line access through a serial connection.

File Writes log entries to a file on the appliance. This file can be archived using the rotate or upload method. The file can be sent as an email.

Depending on the machine type of the appliance, the location of the file can be the local file system or the hard disk array.

SMTP Forwards log entries as an email to configured addresses. The processing rate can be limited.

SOAP Forwards log entries as SOAP messages. The URL can be set. The processing rate can be limited.

syslog Forwards log entries using UDP to a remote syslog daemon. The local address, remote address, remote port, and syslog facility can be set. The processing rate can be limited.

syslog-tcp
Forwards log entries using TCP to a remote syslog daemon. The local

address, remote address, remote port, and syslog facility can be set. An SSL connection to the syslog host can be created. The processing rate can be limited.

Configuring log targets

You can configure a logging target by using the IBM MQ Appliance command line interface.

About this task

Messages in log targets can be restricted by object filters, event category, and event priority. By default, a log target cannot accept messages until it is subscribed to one or more events.

After you have created and configured a logging target, you can add further features by using the global configuration commands that are described in the following topics.

Procedure

To configure a log target:

1. Connect to the IBM MQ Appliance as described in "Command line access" on page 97. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Enter the following command to create your logging target configuration:

   ```
   logging target name
   ```

   Where name specifies the name of the configuration.
4. Use the log target commands to configure your logging target. Use these commands to specify features such as log type, log events, IP address of the target where the log is written. For example, the following commands set up a syslog-tcp target on the remote machine rmach.hursley.ibm.com:

   ```
   mqa(config)# logging target syslog-server
   New Log Target configuration

   mqa(config logging target syslog-server)# summary "Remote logging to rmach.hursley.ibm.com"
   mqa(config logging target syslog-server)# type syslog-tcp
   mqa(config logging target syslog-server)# timestamp syslog
   mqa(config logging target syslog-server)# local-ident "warrior12"
   mqa(config logging target syslog-server)# upload-method ftp
   mqa(config logging target syslog-server)# remote-address "198.51.100.0" "1514"
   mqa(config logging target syslog-server)# local-address 198.51.100.12
   mqa(config logging target syslog-server)# event "all" "notice"
   mqa(config logging target syslog-server)# exit
   mqa(config)# write mem
   ```

Using trace

You can use the **strmqtrc** and **endmqtrc** commands on the command line to start and end tracing.

The **strmqtrc** command has optional parameters to enable you to customize the trace file that is generated. You can trace one or more queue managers. You can trace one or more processes. You can trace specific threads within applications. You

can trace events. You can also specify what level of trace detail you require. For more information about **strmqtrc** and the optional parameters, see “strmqtrc” on page 359.

The **endmqtrc** command has optional parameters to enable you to control which entities the trace is ended for. For more information about these parameters, see “endmqtrc” on page 346.

After you generate the trace files, you can download them from the appliance.

To generate and retrieve trace files from the appliance, complete the following steps:

1. Use the **strmqtrc** command to specify details about the trace information that you want to collect.
2. Use the **endmqtrc** command to end the trace.
3. Use the command **runmqras -section trace** to export the trace information to a file. The command output gives you the file details.
4. Use the **copy** command to download the trace file from the **mqtrace:** URI on the appliance to your local system.

The trace is not formatted on the appliance, but you can format it after you download it, if required, by using the dspmqtrc command, see dspmqrtc in the IBM MQ documentation.

Using trace in the IBM MQ Console

You can trace activity in the IBM MQ Console.

To enable the tracing of the console:

1. Click the help icon in the IBM MQ Console title bar and select **Diagnostics** from the menu.

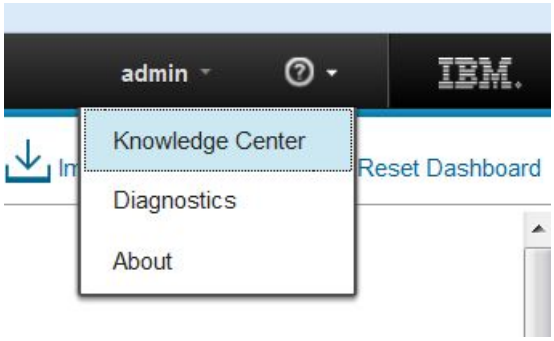

Figure 49. Enabling trace from the IBM MQ Console

2. In the Diagnostics window, select **Enable IBM MQ Console trace**.

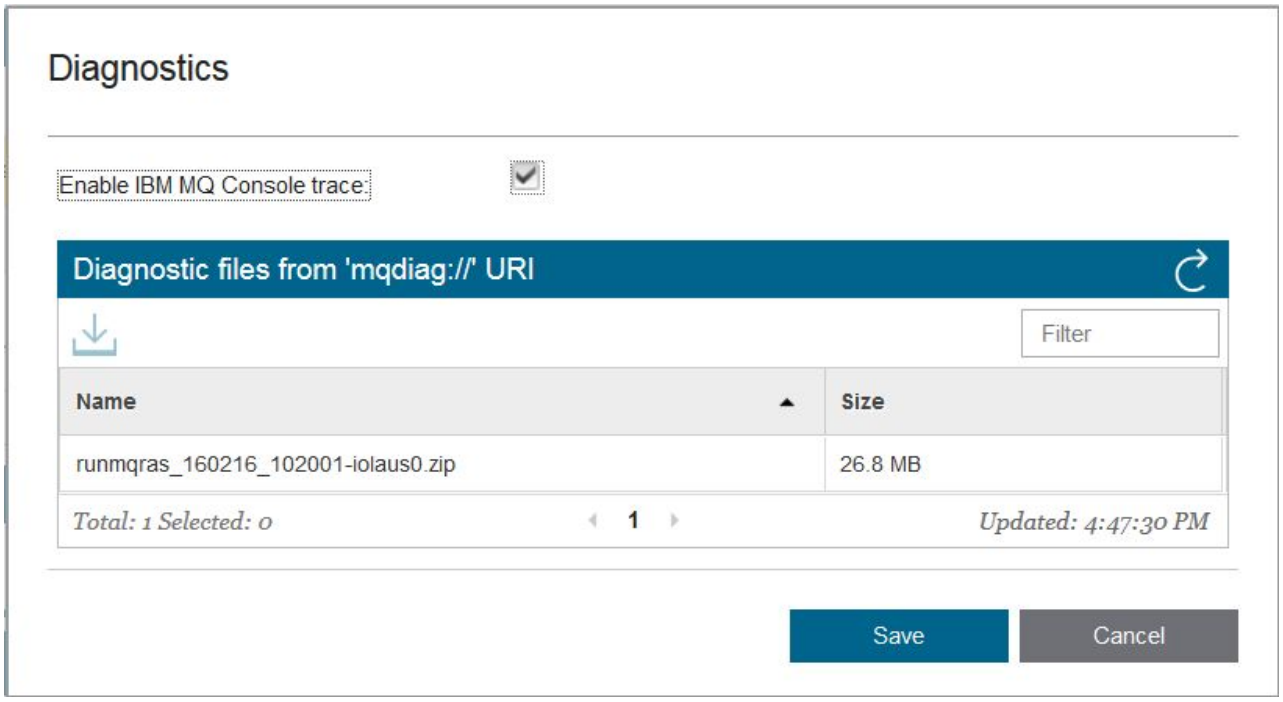

Figure 50. The diagnostics window

3. Click **Save**.

You can then recreate the problem that you are trying to troubleshoot and capture the results in the trace. You then disable trace once more.

1. Click the help icon in the IBM MQ Console title bar and select **Diagnostics** from the menu.
2. In the Diagnostics window, clear **Enable IBM MQ Console trace**.
3. Click **Save**.

Capture save the trace in a file by using the **runmqras -section trace** command from the command line (see "Using trace" on page 316). You can download the trace file from the diagnostics window:

1. Select the file that you want to download.
2. Click the download button .

Resolving a partitioned problem in a high availability configuration

A partitioned problem occurs when the two appliances in a high availability configuration lose the ability to communicate with each other.

If the two appliances in your high availability configuration lose both primary and secondary interface connections, then replication no longer occurs between the two appliances. After the connection is restored, the data replication system detects that there have been independent changes to the same resources on both appliances. This situation is described as a partitioned situation, because the two appliances have two different views of the current state of the queue manager (it is sometimes called a 'split-brain' situation). The queue manager continues to run on the primary, but is stopped on the secondary appliance.

To resolve the situation, you must decide which of the two appliances has the data that you want to retain, you then issue a command that identifies this appliance as the "winner". Data on the other appliance is discarded. The queue manager is then started on the preferred appliance.

To help you decide, you can run the `status` command for the affected queue manager on each appliance. The status command returns status of `partitioned` together with a report of how much out-of-sync data the appliance has for that queue manager.

You identify the winner by running the following commands on the chosen appliance:

1. Connect to the IBM® MQ Appliance as described in Command line access.
2. Log in as a user in the administrators group.
3. Type the following command to enter IBM MQ mode:
   ```
   mqcli
   ```
4. Type the following command on the appliance that you determined to be the "winner":
   ```
   makehaprimary HAQMName
   ```
   Where *HAQMName* is the name of the queue manager. The queue manager then runs on that appliance as it is now the primary.

You can also perform this operation from the IBM MQ Console:

1. Start the IBM MQ Appliance web UI on the appliance that you determined to be the winner and click **MQ Console**.
2. In the queue manager widget, select the queue manager whose data is partitioned and select **More** > **High Availability (HA)** > **Resolve partitioned data**.
3. Confirm that you want the current appliance to become the primary data source for the queue manager.

If the two appliances lose the replication interface, the queue managers are immediately flagged as partitioned. The running queue manager might accumulate out-of-sync data. The other queue manager remains in standby with no out-of-sync data. When the connection is remade, the partitioned state is automatically resolved and replication is resumed.

If your HA queue manager is configured for disaster recovery, and failed over to the recovery appliance when your HA group went out of service, then you might have to resolve data partitioning between the HA group and the recovery appliance. After you have restored your HA group, and resolved data partitioning between the primary and secondary appliances, you must follow the procedure described in "Switching back to the main appliance" on page 217.

Resolving a partitioned problem in a disaster recovery configuration

A partitioned problem is when the queue manager data on one of the appliances in a disaster recovery pair is out of step with the data on the other appliance.

A partitioned problem can arise when the replication link between the two appliances has been lost. It might be the case that a disaster has occurred, and the secondary queue manager has been started on the recovery appliance. When the

main site is restored, the queue manager on the appliance there will be out of step with the queue manager on the recovery appliance.

Depending on how the partitioning occurred, your two appliances could show any of the statuses listed in the following table (this is the status when the previously disconnected connection is restored, but the queue manager is running on the recovery appliance):

Table 31. Partitioned statuses

Main site DR status	Recovery site DR status
Remote appliance(s) unavailable	Partitioned
Partitioned	Remote appliance(s) unavailable
Partitioned	Partitioned

In a partitioned situation you must decide whether to keep the data from the original queue manager, and copy this to the recovery queue manager, or keep the data from the recovery queue manager and copy this to the original queue manager. You use the **makedrprimary** and **makedrsecondary** commands to achieve the desired outcome.

- To keep the data from the queue manager on the recovery appliance:
  1. Ensure the queue managers are stopped.
  2. Specify that the queue manager on the main appliance is the secondary, for example:
     ```
     makedrsecondary -m myqueuemanager
     ```
  3. Specify that the queue manager on the recovery appliance is the primary, for example:
     ```
     makedrprimary -m myqueuemanager
     ```
     Synchronization begins, with the data from the recovery appliance being copied to the main appliance.
  4. When the synchronization is complete, run the **makedrsecondary** command on the queue manager on the recovery appliance, for example:
     ```
     makedrsecondary -m myqueuemanager
     ```
  5. Specify that the queue manager on the main appliance is now the primary, for example:
     ```
     makedrprimary -m myqueuemanager
     ```
  6. Start the queue manager on the main appliance, for example:
     ```
     strmqm myqueuemanager
     ```
- To keep the data from the queue manager on the main appliance:
  1. Ensure the queue managers are stopped
  2. Specify that the queue manager on the recovery appliance is the secondary, for example:
     ```
     makedrsecondary -m myqueuemanager
     ```
  3. Specify that the queue manager on the main appliance is the primary, for example:
     ```
     makedrprimary -m myqueuemanager
     ```
     Synchronization begins, with the data from the main appliance being copied to the recovery appliance.

4. When synchronization is complete, start the queue manager on the main appliance, for example:

```
strmqm myqueuemanager
```

Troubleshooting file copy

If you encounter problems copying files to and from the appliance, try some of the following steps to resolve the problem.

When a copy command fails, make the following checks:

- Ensure that the file name and path that you have specified are correct. The following table gives a list of valid URIs on the appliance.

Table 32. Appliance URIs

URI	Permissions	Description
mqbackup://	copy from copy to	Used for: • user backup and restore • certificate backup and restore • dmpmqcfg output • CCDTs
mqdiag://	copy from	Used for: • runmqras output • amqrfdm • amqspdbg
mqerr://	copy from	Used for: • MQSystem.log and other system logs • FDCs Can also be used to access QM error logs
mqpubcert://	copy from copy to	You can copy certificates to this location. Use the certificate commands (see "Queue manager certificate management commands" on page 374)
mqtemporary://	copy from copy to	This directory is temporary, and does not survive an appliance restart. Under certain severe error conditions that cause other areas of appliance storage to be inaccessible, error reports (FFSTs) are recorded here to assist IBM support staff in problem diagnosis.
mqtrace://	copy from	Used for: • MQ Trace • console trace, errors, and FDCs
mqwebui://	copy from copy to	Saving and loading user dashboards

- Ensure that the appliance is correctly configured with an IP address. You can do this by typing the **show ipaddress** command.
- Ensure that you can ping the server that you are copying a file to or from.
- Ensure that there is an SSH daemon running on the server that you are copying a file to or from.
- Examine the `sshd_config` file on the server that you are copying a file to or from, and ensure that it contains the line `PasswordAuthentication yes`.

If the network connections have not been correctly configured, you might see messages including the information "non-management traffic will be blocked".

This can cause file copy to fail. You can check configurations by using the IBM MQ Appliance web UI. Select **Manage Appliance** > **Network** > **Ethernet Interface**. If any of the interfaces have the status **down**, you must investigate the configuration, and resolve the problem by using one of the following methods:

- Fix the configuration problem, and check that the status changes to **up**.
- Disable the **Block nonmanagement traffic for invalid interface** configuration:
  1. Select **Manage Appliance** > **Network** > **Network Settings**.
  2. Deselect **Block nonmanagement traffic for invalid interface configuration**.

You should also remove cached information for the host that you are copying the file from or to interface on the appliance:

1. At the appliance command line, type `config` to enter configuration mode.
2. Type:
   `no known-host` *`IP_address`*

   Where *IP_address* is the IP address of the host that you are copying the file from or to.

Help with using `runmqras`

Tips on using the **runmqras** command to collect troubleshooting information.

You should use **runmqras** command only when instructed by IBM support. When so instructed, proceed as follows:

1. From the command line, run the **runmqras** as instructed by IBM support (see "runmqras" on page 351 for more information about the command).
2. Retrieve the file that **runmqras** created from the appliance. The file is located under the `mqdiag://` URI, and has a name of the form `runmqras_`*`timestamp`*`.zip`.You can retrieve the file by using the appliance command line, or by using the IBM MQ Console.
   - From the command line, enter the following commands:
     ```
     config
     copy mqdiag://runmqras_YYMMDD_HHmmss.zip scp://user@host//home/user
     ```
   - In the IBM MQ Console, use the File Manager to navigate to the `mqdiag://` URI and save the `runmqras_`*`YYMMDD_HHmmss`*`.zip` file to your local computer.
3. Send the file to IBM Support.

You should take note of the following features when you use the **runmqras** command.

- Use of the `runmqras -section all` command is likely to fill up all allocated space. This is especially true if there are multiple default-sized queue managers.

- Use `runmqras -section trace` to collect the trace files.
- If the Web UI output (`wlp_dump.zip`) is not picked up by **runmqras**, it could be because the dumping command timed out. Try running `runmqras -section webui` again.

Chapter 11. Reference

You can use the reference information in this section to accomplish the tasks that address your business needs.

IBM MQ commands

Use the IBM MQ commands to work with messaging features.

Using commands

To use the IBM MQ commands, connect to the appliance as described in "Command line access" on page 97.

You must enter MQ administration mode before issuing commands. To enter MQ administration mode, type `mqcli` on the command line. You see the prompt mqa(mqcli)#.

Command help

The supported IBM MQ commands can be viewed on the command line. These commands are divided into categories, including administration commands, diagnosis commands, user commands, and certificate commands. To view a list of the available categories, enter the following command from the IBM MQ administration mode:

```
help
```

To view the commands that are in a specific category, enter the following command from the IBM MQ administration mode:

```
help category
```

To view detailed help about a particular command, enter one of the following commands from the IBM MQ administration mode:

- `help commandName`
- `? commandName`

where *commandName* is the name of the command that you want to view the help for.

The command descriptions in the following topics use railroad diagrams for command syntax. for an explanation of how to use these diagrams, see How to read railroad diagrams in the IBM MQ documentation.

IBM MQ Control commands

You can use the IBM MQ control commands to manage queue managers and to perform various utility functions, such as running MQSC commands.

The IBM MQ control commands can be run from the command line interface in MQ command mode. To enter MQ command mode, type `mqcli`.

The commands, including the command name itself, the flags, and any arguments, are case-sensitive. For example:

© Copyright IBM Corp. 2015, 2016

```
crtmqm -u SYSTEM.DEAD.LETTER.QUEUE QM1
```

- The command name must be `crtmqm`, not `CRTMQM`
- The flag must be `-u`, not `-U`
- The dead-letter queue is called SYSTEM.DEAD.LETTER.QUEUE
- The argument is specified as `QM1`, which is different from `qm1`

crtmqm

Create a queue manager.

Purpose

You can use the **crtmqm** command to create a queue manager.

Syntax

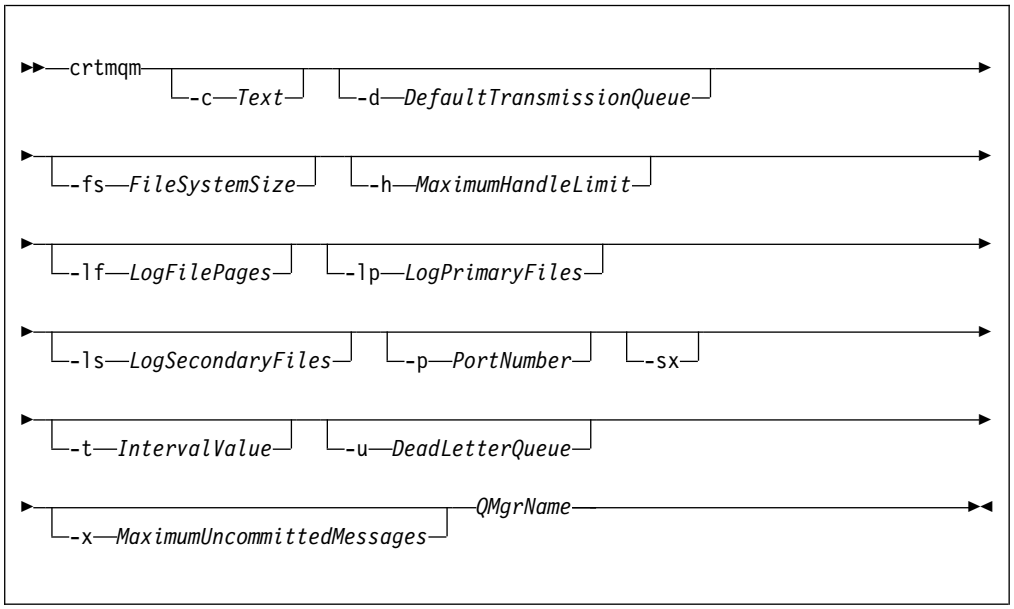

Parameters

QMgrName
Specifies the name of the queue manager that you want to create.

The queue manager name must be the last parameter that is specified in the command.

The name can contain up to 48 characters. The following characters can be used:
`0-9 A-Z a-z . / _ %`

The name of the queue manager must be unique on the IBM MQ Appliance. If the queue manager connects to other queue managers, the queue manager names must be unique within that group of queue managers.

This parameter is required.

-c *Text*
Specifies descriptive text for this queue manager.

You can use up to 64 characters. If you include special characters, enclose the description in single quotation marks. The maximum number of characters is reduced if you are using a double-byte character set (DBCS).

The default value is all blanks.

-d *DefaultTransmissionQueue*
Specifies the name of the local transmission queue where remote messages are put if a transmission queue is not explicitly defined for their destination.

There is no default value.

-fs *FileSystemSize*
Specifies that the queue manager is created with the file system size *FileSystemSize*. If you do not specify this argument, the file system size defaults to 64 GB.

FileSystemSize is a numeric value, which is specified in GB. You can specify a value in MB by entering the value followed by the character `M`. For example, to specify a *FileSystemSize* of 3 GB, enter `3`. To specify a *FileSystemSize* of 1024 MB, enter `1024M`.

For the appliance the minimum value is 128 MB.

Note: After a queue manager is created you cannot resize the file system; ensure the value that is specified here is sufficient for the current and any future workload.

-h *MaximumHandleLimit*
Specifies the maximum number of handles that an application can open at the same time.

Specify a value in the range 1 - 999999999.

The default value is 256.

-lf *LogFilePages*
Specifies the number of log file pages to use for the log files.

The log data is held in a series of files called log files. The log file size is specified in units of 4 KB pages.

The default number of log file pages is 4096, giving a log file size of 16 MB. The minimum number of log file pages is 64 and the maximum is 65535.

-lp *LogPrimaryFiles*
Specifies the log files that are allocated when the queue manager is created.

The minimum number of primary log files you can have is 2 and the maximum is 510. The default is 3. The total number of primary and secondary log files must not exceed 511 and must not be less than 3.

You can change this value after the queue manager is created. However, the change is not effective until the queue manager is restarted.

-ls *LogSecondaryFiles*
Specifies the log files that are allocated when the primary files are exhausted.

The minimum number of secondary log files you can have is 2 and the maximum is 509. The default is 2. The total number of primary and secondary log files must not exceed 511 and must not be less than 3.

You can change this value after the queue manager is created. However, the change is not effective until the queue manager is restarted.

-p *PortNumber*
Create a managed TCP listener on the specified port.

Specify a valid port value to create a TCP listener object that uses the specified port. The new listener is called SYSTEM.LISTENER.TCP.1. This listener is under queue manager control, and is started and stopped along with the queue manager.

-sx
Specifies that the queue manager is a high availability (HA) queue manager.

The queue manager starts automatically as part of the HA group.

-t *IntervalValue*
Specifies the trigger time interval in milliseconds for all queues that are controlled by this queue manager.

That is, after the queue manager receives a trigger-generating message, triggering is suspended for the length of time that is specified by *IntervalValue*.

Specify a value in the range 0 - 999999999.

The default value is 999999999 milliseconds. This value effectively means that triggering is disabled after the first trigger message.

-u *DeadLetterQueue*
Specifies the name of the local queue that is to be used as the dead-letter (undelivered-message) queue.

The default is no dead-letter queue.

-x *MaximumUncommittedMessages*
Specifies the maximum number of uncommitted messages under any one sync point.

The uncommitted messages are the sum of the following messages:

- The number of messages that can be retrieved from queues
- The number of messages that can be put on queues
- Any trigger messages that are generated within this unit of work

The limit that is specified does not apply to messages that are retrieved or put outside a sync point.

Specify a value in the range 1 - 999999999.

The default value is 10000 uncommitted messages.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- After you create the queue manager, you can use the **strmqm** command to start the queue manager. A high availability queue manager is started automatically after creation, so you do not need to start it by using **strmqm**.
- When a queue manager is created, the default and system objects are also created. These objects are listed in System and default objects in the IBM MQ documentation.
- For more information about this command in IBM MQ, see ctrmqm in the IBM MQ documentation.

Examples

- The following command creates a queue manager that is called QM1, with a description of example queue manager, and creates the system and default objects:

  ```
  crtmqm -c "example queue manager" QM1
  ```

- The following command creates a queue manager that is called QM2. It creates the system and default objects, sets the trigger interval to 5000 milliseconds (5 seconds), and specifies SYSTEM.DEAD.LETTER.QUEUE as its dead-letter queue.

  ```
  crtmqm -t 5000 -u SYSTEM.DEAD.LETTER.QUEUE QM2
  ```

Related commands

- strmqm (Start queue manager)
- endmqm (End queue manager)
- dltmqm (Delete queue manager)

dltmqm

Delete a queue manager.

Purpose

You can use the **dltmqm** command to delete a queue manager.

Syntax

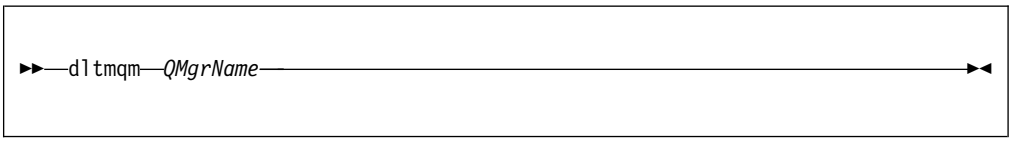

Parameters

QMgrName

Specifies the name of the queue manager that you want to delete.

This parameter is required.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter mqcli on the command line. To exit the IBM MQ administration mode, enter exit on the command line.
- Before you delete the queue manager, you must end the queue manager by using the **endmqm** command.
- For more information about this command in IBM MQ, see dltmqm in the IBM MQ documentation.

Examples

- The following command deletes the queue manager QM1.

  ```
  dltmqm QM1
  ```

Related commands

- crtmqm (Create queue manager)
- strmqm (Start queue manager)
- endmqm (Delete queue manager)

dmpmqcfg

Dump the configuration of a queue manager.

Purpose

You can use the **dmpmqcfg** command to dump the configuration of a queue manager.

Syntax

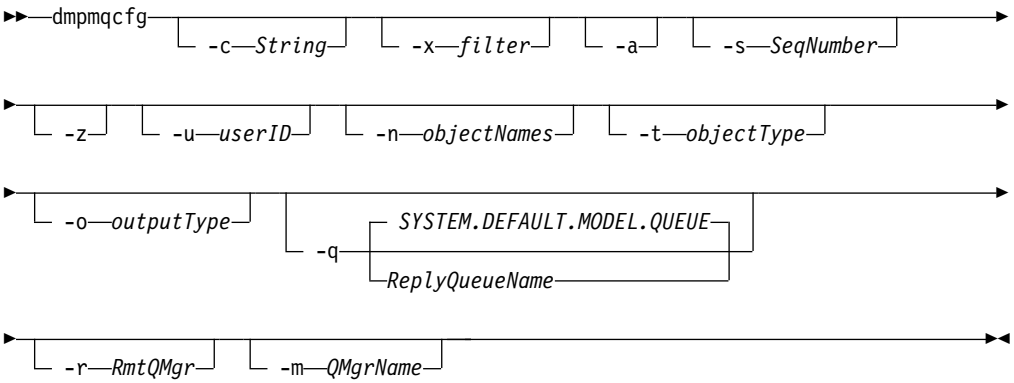

Parameters

-c *string*

Specifies that a client mode connection is used to connect to the queue manager.

string can take one of the following values:

default

Specifies that the default client connection process is used.

"DEFINE CHANNEL(*chlname*) CHLTYPE(CLNTCONN) CONNAME('*connname*')"

Specifies that the specific client channel specified by *chlname* is used to connect to the queue manager at *connname*.

connname specifies the location of the queue manager in the following format *host*(*portnumber*)

If **-c** is omitted, the command connects to the queue manager by using server bindings. If that connection fails, client bindings are used.

-x *filter*

Specifies that the procedure is filtered.

filter can be one of the following values:

object

authority records

channel authentication

subscriptions

all

The default value is `all`.

-a Specifies that object definitions show all attributes.

The default is to return only attributes that differ from the defaults for the object type.

-s *SeqNumber*
Specifies that the channel sequence number for sender, server, and cluster sender channel types is reset to the value specified.

SeqNumber must be in the range 1 - 999999999.

-z Specifies that the command runs in silent mode.

All warnings, such as those that appear when attributes from a queue manager of a higher command level are inquired, are suppressed.

-n *objectNames*
Specifies that the definitions produced by object or profile name are filtered.

The object or profile name can contain a single asterisk. The `*` option can be placed only at the end of the entered filter string.

-t *objectType*
Specifies a single type of object to export.

objectType can be one of the following values:

all All object types.

authinfo
An authentication information object.

channel
A channel (including MQTT channel type).

comminfo
A communications information object.

listener
A listener.

namelist
A namelist.

process
A process.

queue A queue.

qmgr A queue manager.

service
A service.

topic A topic.

The default value is `all`.

-o *outputType*
Specifies the type of output for the command.

outputType can be one of the following values:

mqsc Multi-line MQSC that can be used as direct input to **runmqsc**

1line MQSC with all attributes on a single line for line diffing

setmqaut
setmqaut statements; valid only when `-x authrec` is specified

grtmqaut
Generates IBM i syntax for granting access to the objects.

The default value is `mqsc`.

-u *userID*
If a userID is specified, a password is requested.

-q Specifies the name of the reply-to queue used when configuration information is retrieved.

-r Specifies the name of the remote queue manager/transmit queue when queued mode is used.

If this parameter is omitted, the configuration for the directly connected queue manager (specified with the **-m** parameter) is dumped.

-m Specifies the name of the queue manager to connect to.

The default value is the default queue manager.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- If any object is not at the default value, the `-a` option must be used if the dumped configuration is used to restore the configuration.
- The **dmpmqcfg** command dumps only subscriptions of type MQSUBTYPE_ADMIN, that is, only subscriptions that are created by using the MQSC command **DEFINE SUB** or its PCF equivalent. The output from **dmpmqcfg** is a runmqsc command to enable the administration subscription to be re-created. Subscriptions that are created by applications by using the MQSUB MQI call of type MQSUBTYPE_API are not part of the queue manager configuration, even if durable, and so are not dumped by **dmpmqcfg**.
- The user must have MQZAO_OUTPUT (+put) authority to access the command input queue (SYSTEM.ADMIN.COMMAND.QUEUE) and MQZAO_DISPLAY (+dsp) authority to access the default model queue (SYSTEM.DEFAULT.MODEL.QUEUE), to be able to create a temporary dynamic queue if the default reply queue is used. The user must also have MQZAO_CONNECT (+connect) and MQZAO_INQUIRE (+inq) authority for the queue manager, and MQZAO_DISPLAY (+dsp) authority for every object that is requested.
- For more information about this command in IBM MQ, see dmpmqcfg in the IBM MQ documentation.

Examples

- The following command dumps the queue manager configuration for a queue manager `QM1`:

```
dmpmqcfg -m QM1
```

dspmq

Display information about queue managers.

Purpose

You can use the **dspmq** command to display the names and details of the queue managers on the IBM MQ Appliance.

Syntax

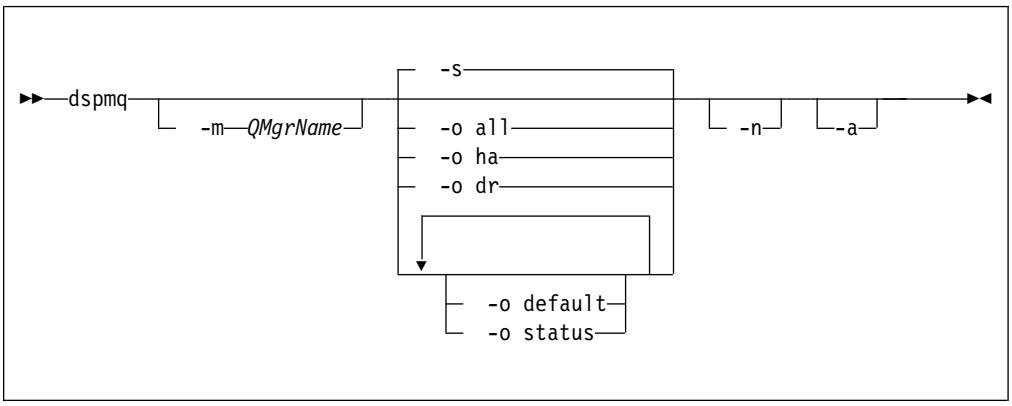

Parameters

-a Specifies that information about only the active queue managers is displayed.

A queue manager is active one or more of the following statements are true:

- The queue manager is running
- A listener for the queue manager is running
- A process is connected to the queue manager

-m *QMgrName*

Specifies which queue manager to display the details for.

If no queue manager name is specified, all queue managers are displayed.

-n Specifies that the translation of output strings is suppressed.

-s

Specifies that the operational status of the queue managers is displayed.

This parameter is the default status setting. It is equivalent to **-o status**.

-o all

Specifies that the operational status of the queue managers is displayed.

-o default

Specifies that the default queue manager status is displayed.

-o ha

Specifies that the HA type is displayed.

-o dr

Specifies that disaster recovery information is displayed. Displays the port that the data replication listener on both appliances uses and the IP address used by the remote appliance.

-o status

Specifies that the operational status of the queue managers is displayed.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- The queue manager can be in any of the following states:
  - Starting
  - Running
  - Running as standby
  - Running elsewhere
  - Quiescing
  - Ending immediately
  - Ending pre-emptively
  - Ended normally
  - Ended immediately
  - Ended unexpectedly
  - Ended pre-emptively
  - Status not available
- For more information about this command in IBM MQ, see dspmq in the IBM MQ documentation.

Examples

- The following command displays queue managers on the appliance:

```
dspmq -o all
```

dspmqrte

Determine the route that a message has taken through a queue manager network.

Purpose

You can use the **dspmqrte** to generate a trace-route message and put it into a queue manager network. As the trace-route message travels through the queue manager network, activity information is recorded. When the trace-route message reaches its target queue, the activity information is collected and displayed.

Syntax

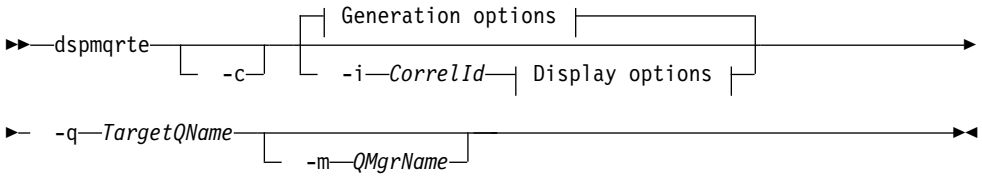

Generation options:

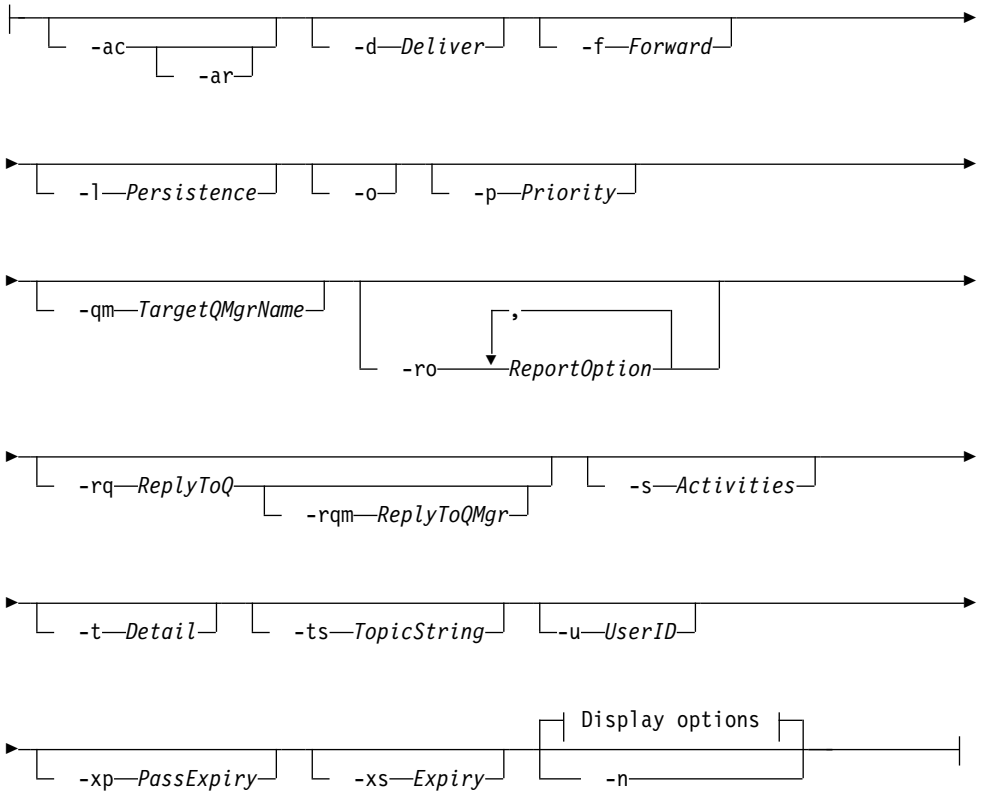

Display options:

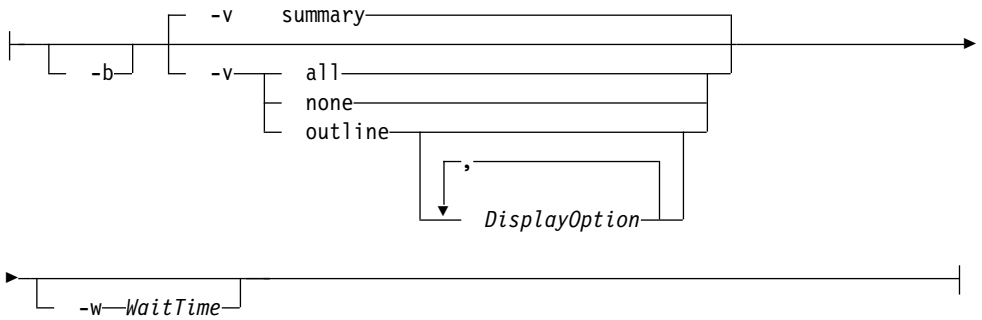

Parameters

-q *TargetQName*
Specifies the name of the target queue to send the trace-route message to.

If the command is being used to view previously gathered activity information, *TargetQName* specifies the name of the queue where the activity information is stored.

This parameter is required.

-c
Specifies that the command connects as a client application.

-i *CorrelId*
Specifies the message identifier of the original trace-route message when displaying previously accumulated activity information.

There can be many activity reports and trace-route reply messages on the queue specified by **-q TargetQName**. Therefore, **-i** can be used to identify the activity reports, or a trace-route reply message, related to a specific trace-route message.

Specify *CorrelId* as a 48 character hexadecimal string.

-m *QMgrName*
Specifies the name of the queue manager to which the command connects.

The name can contain up to 48 characters.

The default value is the default queue manager.

The following parameters are used when the command is used to put a trace-route message into a queue manager network. That is, the parameters are the generation options:

-ac
Specifies that activity information is to be accumulated within the trace-route message.

If you do not specify this parameter, activity information is not accumulated within the trace-route message.

-ar
Specifies that a trace-route reply message containing all accumulated activity information is generated in the following circumstances:

- The trace-route message is discarded by a queue manager.
- The trace-route message is put to a local queue (target queue or dead-letter queue) by a queue manager.
- The number of activities performed on the trace-route message exceeds the value of specified in *-s Activities*.

If you do not specify this parameter, a trace-route reply message is not requested.

-d *Deliver*
Specifies whether the trace-route message is to be delivered to the target queue on arrival.

Deliver can be one of the following values:

yes On arrival, the trace-route message is put to the target queue, even if the queue manager does not support trace-route messaging

no On arrival, the trace-route message is not put to the target queue.

The default value is **no**.

-f *Forward*
Specifies the type of queue manager that the trace-route message can be forwarded to.

Forward can be one of the following values:

all The trace-route message is forwarded to any queue manager.

If forwarded to a queue manager before Version 6.0, the trace-route message is not recognized and can be delivered to a local queue despite the value of the `-d` parameter.

supported
The trace-route message is only forwarded to a queue manager that honors the value of the `-d` parameter.

The default value is `supported`.

-l *Persistence*
Specifies the persistence of the generated trace-route message.

Persistence can be one of the following values:

yes The generated trace-route message is persistent. (MQPER_PERSISTENT)

If you use this value, you must specify the parameter **-rq ReplyToQ**. The reply-to queue must not resolve to a temporary dynamic queue.

no The generated trace-route message is not persistent. (MQPER_NOT_PERSISTENT).

q The generated trace-route message inherits its persistence value from the queue specified by **-q TargetQName**. (MQPER_PERSISTENCE_AS_Q_DEF).

A trace-route reply message, or any report messages, returned shares the same persistence value as the original trace-route message.
The default value is `no`.

-o Specifies that the target queue is not bound to a specific destination.

Typically this parameter is used when the trace-route message is to be put across a cluster. The target queue is opened with option MQOO_BIND_NOT_FIXED.

If you do not specify this parameter, the target queue is bound to a specific destination.

-p *Priority*
Specifies the priority of the trace-route message.

The value of *Priority* is either greater than or equal to 0, or MQPRI_PRIORITY_AS_Q_DEF. MQPRI_PRIORITY_AS_Q_DEF specifies that the priority value is taken from the queue specified by **-q TargetQName**.
The default is that the priority value is taken from the queue specified by **-q TargetQName**.

-qm *TargetQMgrName*
Specifies the target queue manager for the target queue.

The target queue is specified with **-q TargetQName**.

The default is that the queue manager to which the command is connected is used as the reply-to queue manager.

-ro *ReportOption*
ReportOption can be one or more of the following values specified in a comma-separated list:

none Specifies that no report options are set.

activity
The report option MQRO_ACTIVITY is set.

coa The report option MQRO_COA_WITH_FULL_DATA is set.

cod The report option MQRO_COD_WITH_FULL_DATA is set.

exception
The report option MQRO_EXCEPTION_WITH_FULL_DATA is set.

expiration
The report option MQRO_EXPIRATION_WITH_FULL_DATA is set.

discard
The report option MQRO_DISCARD_MSG is set.

The default value is `activity, discard`.

-rq *ReplyToQ*
Specifies the name of the reply-to queue that all responses to the trace-route message are sent to.

If the trace-route message is persistent, or if the `-n` parameter is specified, a reply-to queue must be specified that is not a temporary dynamic queue.

If you do not specify this parameter, the system default model queue, SYSTEM.DEFAULT.MODEL.QUEUE is used as the reply-to queue. Using this model queue causes a temporary dynamic queue to be created.

-rqm *ReplyToQMgr*
Specifies the name of the queue manager where the reply-to queue is located.

The name can contain up to 48 characters.

If you do not specify this parameter, the queue manager to which the command is connected is used as the reply-to queue manager.

-s *Activities*
Specifies the maximum number of recorded activities that can be performed on behalf of the trace-route message before it is discarded.

This parameter prevents the trace-route message from being forwarded indefinitely if caught in an infinite loop.

The value of *Activities* is either greater than or equal to 1, or MQROUTE_UNLIMITED_ACTIVITIES. MQROUTE_UNLIMITED_ACTIVITIES specifies that an unlimited number of activities can be performed on behalf of the trace-route message.

If you do not specify this parameter, an unlimited number of activities can be performed on behalf of the trace-route message.

-t *Detail*
Specifies the activities that are recorded.

Detail can be one of the following values:

low Activities performed by user-defined application are recorded only.

medium
Activities specified in **low** are recorded. Additionally, activities performed by MCAs are recorded.

high Activities specified in **low**, and **medium** are recorded. MCAs do not expose any further activity information at this level of detail. This option is available to user-defined applications that are to expose further activity information only. For example, if a user-defined application determines the route a message takes by considering certain message characteristics, the routing logic can be included with this level of detail.

The default value is `medium`.

-ts *TopicString*
Specifies a topic string to which the command is to publish a trace-route message, and puts the command into topic mode.

In this mode, the command traces all of the messages that result from the publish request.

-u *userID*
User ID to use for connecting to a queue manager.

-xp *PassExpiry*
Specifies whether the report option MQRO_DISCARD_MSG and the remaining expiry time from the trace-route message is passed on to the trace-route reply message.

PassExpiry can be one the following values:

yes The report option MQRO_PASS_DISCARD_AND_EXPIRY is specified in the message descriptor of the trace-route message.

If a trace-route reply message, or activity reports, are generated for the trace-route message, the MQRO_DISCARD_MSG report option (if specified), and the remaining expiry time are passed on.

no The report option MQRO_PASS_DISCARD_AND_EXPIRY is not specified.

If a trace-route reply message is generated for the trace-route message, the discard option and remaining expiry time from the trace-route message are not passed on.

The default value is `yes`.

-xs *Expiry*
Specifies the expiry time for the trace-route message, in seconds.

The default value is 60.

-n Specifies that activity information returned for the trace-route message is not to be displayed.

If this parameter is accompanied by a request for a trace-route reply message (**-ar**), or any of the report generating options (**-ro ReportOption**), then a specific (non-model) reply-to queue must be specified using **-rq ReplyToQ**.

After the trace-route message is put to the specified target queue, a 48 character hexadecimal string is returned containing the message identifier of the trace-route message. The message identifier can be used by the command to display the activity information for the trace-route message at a later time. This can be done using the **-i CorrelId** parameter.

By default, activity report messages are requested.

The following parameters are used when the command is used to display collected activity information. That is, the parameters are the display options:

-b Specifies that the command only browses activity reports or a trace-route reply message related to a message.

This parameter allows activity information to be displayed again at a later time.

If you do not specify this parameter, the command gets activity reports or a trace-route reply message related to a message, and deletes them.

-v summary | all | none | outline *DisplayOption*
The values can be the following values:

summary
The queues that the trace-route message was routed through are displayed.

all All available information is displayed.

none No information is displayed.

outline ***DisplayOption***
Specifies display options for the trace-route message.

DisplayOption can be one or more of the following values, using a comma as a separator:

activity
All non-PCF group parameters in Activity PCF groups are displayed.

identifiers
Values with parameter identifiers MQBACF_MSG_ID or MQBACF_CORREL_ID are displayed.

This value overrides `msgdelta`.

message
All non-PCF group parameters in Message PCF groups are displayed.

When this value is specified, you cannot specify `msgdelta`.

msgdelta
All non-PCF group parameters in Message PCF groups, that have changed since the last operation, are displayed.

When this value is specified, you cannot specify `message`.

operation
All non-PCF group parameters in Operation PCF groups are displayed.

traceroute
All non-PCF group parameters in TraceRoute PCF groups are displayed.

If no values are supplied for *DisplayOption* the application name, the type of each operation, and any operation specific parameters are displayed.

The default value is `summary`.

-w *WaitTime*
Specifies the time, in seconds, that the command waits for activity reports, or a trace-route reply message, to return to the specified reply-to queue.

The default value is the expiry time of the trace-route message, plus 60 seconds.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- For more information about this command in IBM MQ, see dspmqrte in the IBM MQ documentation.

Examples

- The following command puts a trace-route message into a queue manager network with the target queue specified as `TARGET.Q`. Providing queue managers on route are enabled for activity recording, activity reports are generated. Depending on the queue manager attribute, ACTIVREC, activity reports are either delivered to the reply-to queue `ACT.REPORT.REPLY.Q`, or are delivered to a system queue. The trace-route message is discarded on arrival at the target queue.

  ```
  dspmqrte -q TARGET.Q -rq ACT.REPORT.REPLY.Q
  ```

  Providing one or more activity reports are delivered to the reply-to queue, `ACT.REPORT.REPLY.Q`, the command orders and displays the activity information.
- The following command puts a trace-route message into a queue manager network with the target queue specified as `TARGET.Q`. Activity information is accumulated within the trace-route message, but activity reports are not generated. On arrival at the target queue, the trace-route message is discarded. Depending on the value of the target queue manager attribute, ROUTEREC, a trace-route reply message can be generated and delivered to either the reply-to queue, `TRR.REPLY.TO.Q`, or to a system queue.

  ```
  dspmqrte -ac -ar -ro discard -rq TRR.REPLY.TO.Q -q TARGET.Q
  ```

  Providing a trace-route reply message is generated, and delivered to the reply-to queue `TRR.REPLY.TO.Q`, the command orders and displays the activity information that was accumulated in the trace-route message.

dspmqtrn

Display in-doubt and heuristically completed transactions.

Purpose

You can use the **dspmqtrn** command to display details of transactions. The transactions that can be displayed include transactions that are coordinated by both the IBM MQ Appliance and external transaction managers.

Syntax

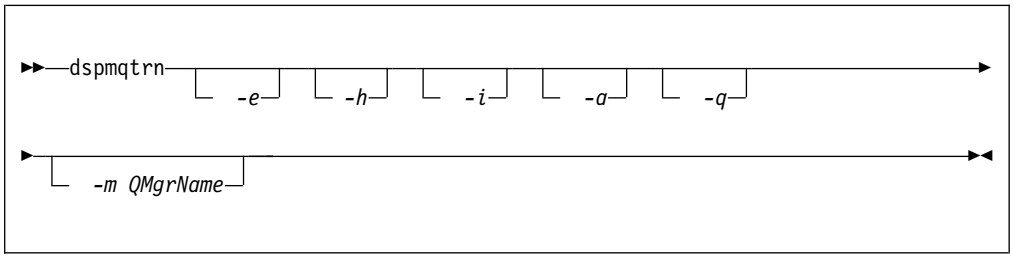

Parameters

-e Specifies that details of externally coordinated, in-doubt transactions are displayed.

These transactions are transactions for which the IBM MQ Appliance has been asked to prepare to commit, but has not yet been informed of the transaction outcome.

-h Specifies that details of externally coordinated, heuristically completed transactions are displayed.

These transactions are transactions that are resolved by the **rsvmqtrn** command, but that the external transaction coordinator has yet to acknowledge with an **xa-forget** command.

-i Specifies that details of internally coordinated, in-doubt transactions are displayed.

These transactions are transactions for which each resource manager has been asked to prepare to commit, but the IBM MQ Appliance has yet to inform the resource managers of the transaction outcome.

-a Specifies that a list of all transactions known to the queue manager are displayed.

The returned data includes transaction details for all transactions known to the queue manager. If a transaction is currently associated with an IBM MQ Appliance application connection, information related to that application connection is also returned.

-q Specifies that all the data from the `-a` parameter and a list of up to 100 unique objects updated within the transaction are displayed.

If more than 100 objects are updated in the same transaction, only the first 100 distinct objects are listed for each transaction.

Specifying this parameter on its own is the same as specifying `-a -q`.

-m *QMgrName*
Specifies the name of the queue manager for which to display transactions.

The default value is the default queue manager.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- If you do not specify `-e`, `-h`, or `-i`, details of both internally and externally coordinated in-doubt transactions are displayed. Details of externally coordinated, heuristically completed transactions are not displayed.
- Not all of the fields are appropriate for all transactions. When the fields are not meaningful, they are displayed as blank.
- For more information about this command in IBM MQ, see dspmqtrn in the IBM MQ documentation.

Examples

- The following command shows externally coordinated, in-doubt transactions for the queue manager `QM1`:

```
dspmqtrn -e -m QM1
```

Related commands

- rsvmqtrn

dspmqver

Display version and build information.

Purpose

You can use the **dspmqver** command to display version and build information for the IBM MQ Appliance.

Syntax

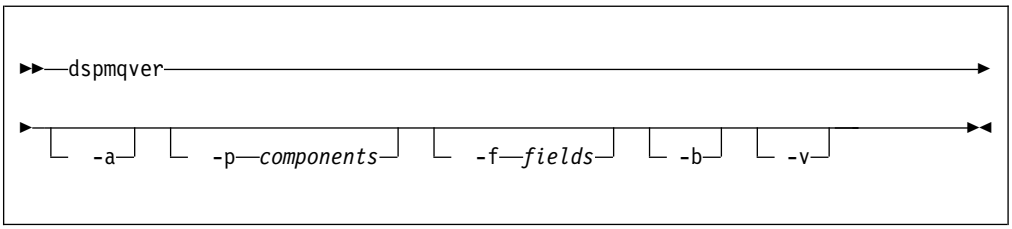

Parameters

-a Specifies that information about all fields and components is displayed.

-p *Components*
Specifies that only information about the components that are specified is displayed.

Multiple components can be specified as a sum of the values of the required components. The components have the following values:

1 IBM MQ Appliance

64 GSKit

128 IBM MQ Advanced Message Security

The default value is 1.

-f *Fields*
Specifies that only information about the fields that are specified is displayed.

Multiple components can be specified as a sum of the values of the required fields. The fields have the following values:

1 Name

2 Version, in the form `V.R.M.F`: Where `V`=Version, `R`=Release, `M`=Modification, and `F`=Fix pack

4 Level

8 Build type

The default value is 15.

-b Specifies that header information is omitted.

-v Specifies that verbose output is displayed.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- For more information about this command in IBM MQ, see dspmqver in the IBM MQ documentation.

Examples

- The following command displays verbose version and build information:

```
dspmqver -v
```

endmqm

Stop a queue manager.

Purpose

You can use the **endmqm** command to stop a queue manager. This command stops a queue manager in one of three modes:

- Controlled or quiesced shutdown
- Immediate shutdown
- Preemptive shutdown

Syntax

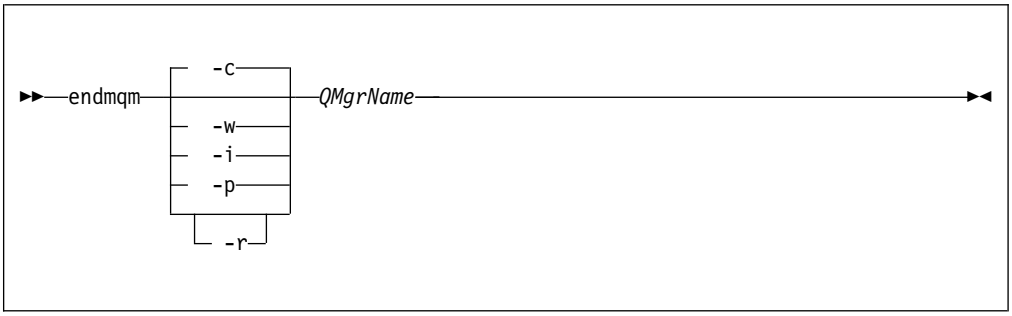

Parameters

QMgrName
Specifies the name of the message queue manager that you want to stop.

This parameter is required.

-c Specifies that the queue manager ends in a controlled (or quiesced) shutdown.

In a controlled shutdown, the queue manager stops after all applications are disconnected. Any MQI calls currently being processed are completed. Control is returned to you immediately and you are not notified of when the queue manager is stopped.

This parameter is the default.

-i Specifies that the queue manager ends in an immediate shutdown.

In an immediate shutdown, the queue manager stops after it all the MQI calls currently being processed are completed. Any MQI requests made after the command starts fail. Any incomplete units of work are rolled back when the queue manager is next started. Control is returned after the queue manager ends.

-p Specifies that the queue manager ends in a preemptive shutdown.

In a preemptive shutdown, the queue manager might stop without waiting for applications to disconnect or for MQI calls to complete. This behavior can give unpredictable results for your applications. Therefore, use this type of shutdown only after other **endmqm** commands fail to stop the queue manager.

-r Specifies that client connectivity can be re-established with other queue managers in their queue manager group.

The client might not reconnect to the same queue manager. Depending on the MQCONNX reconnect option the client uses, and the definition of the queue manager group in the client connection table, the client might reconnect to a different queue manager. You can configure the client to force it to reconnect to the same queue manager.

-w Specifies that the queue manager ends in a wait shutdown.

In a wait shutdown, the queue manager stops after all applications are disconnected. Any MQI calls currently being processed are completed. Control is returned to you after the queue manager stops.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes

mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

- This command does not affect the attributes of the queue manager.
- The **endmqm** command affects any client application that is connected to the queue manager by a server-connection channel. The effect is equivalent to a STOP CHANNEL command in one of the following modes:
  - If the `-c`, or `-w` parameters are used, the mode is QUIESCE.
  - If the `-i` parameter is used, the mode is FORCE.
  - If the `-p` parameter is used, the mode is TERMINATE.
- If an **dspmq** command is entered in the time between the applications disconnecting and the queue manager stopping, the **dspmq** command might report the status as `Ending immediately`, even if a controlled shutdown was requested.
- For more information about this command in IBM MQ, see endmqm in the IBM MQ documentation.

Examples

- The following command ends the queue manager that is named `QM1` in a controlled way:

  ```
  endmqm QM1
  ```

- The following command ends the queue manager that is named `QM2` immediately:

  ```
  endmqm -i QM2
  ```

Related commands

- crtmqm (Create queue manager)
- strmqm (Start queue manager)
- dltmqm (Delete queue manager)

endmqtrc

End trace for some or all of the entities that are being traced.

Purpose

You can use the **endmqtrc** command to end tracing for a specified entity, or for all entities.

Syntax

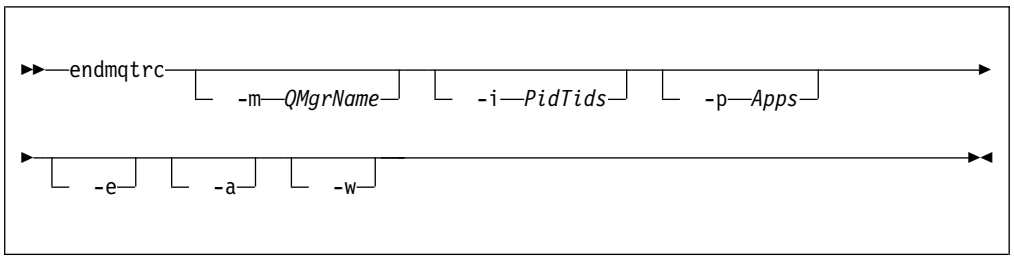

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which to end tracing.

The *QMgrName* supplied must match exactly the *QMgrName* supplied on the **strmqtrc** command. If the **strmqtrc** command used wildcard characters, the **endmqtrc** command must use the same wildcard characters.

A maximum of one `-m` flag can be supplied on the command.

-i *PidTids*
Specifies the process identifier (PID) and thread identifier (TID) for which to end tracing.

You cannot use the `-i` flag with the `-e` flag.

This parameter must be used only under the guidance of IBM Service personnel.

-p *Apps*
Specifies the processes for which to end tracing.

Specify *Apps* as a comma-separated list, with each name in the list specified exactly as the program name would be displayed in the "Program Name" FDC header. You can use an asterisk (*) as a wildcard to match zero or more characters. You can use a question mark (?) to match a single character.

You cannot use the `-p` flag with the `-e` flag.

-e Specifies that early tracing of all processes ends.

You cannot use the `-e` flag with the `-m` flag, the `-i` flag, or the `-p` flag.

-a Ends all tracing.

This flag must be specified alone.

-w Restrict triggering of trace to applications run by an IBM MQ Appliance administrator.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Specifying **endmqtrc** with no parameters has the same effect as specifying **endmqtrc -e**.
- For more information about this command in IBM MQ, see endmqtrc in the IBM MQ documentation.

Examples

- The following command ends tracing of data for a queue manager called `QM1`:

```
endmqtrc -m QM1
```

- The following command ends tracing for queue manager `QM2` only. Any other traces that are running are not affected:

```
endmqtrc -m QM2
```

Related commands

- “strmqtrc” on page 359

mqrc

Display information about return codes.

Purpose

You can use the **mqrc** command to display information about symbols, return codes, and AMQ messages. You can specify a range of return codes or AMQ messages, or you can specify specific return codes or AMQ messages.

Syntax

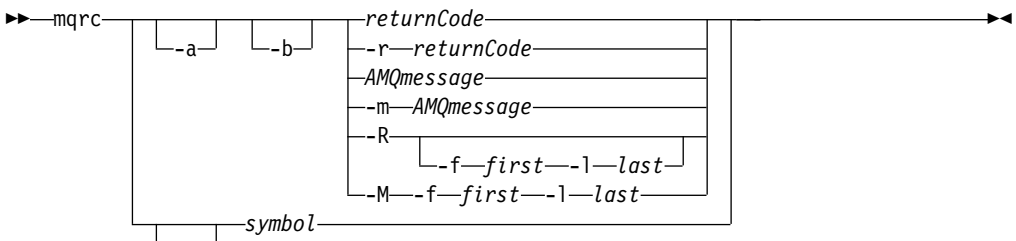

Parameters

returnCode
Specifies the return code to display.

AMQmessage
Specifies the AMQ message to display.

symbol
Specifies the symbol to display.

-a Specifies that all severities are tried to find message text.

-b Specifies that messages are displayed without extended information.

-m *AMQmessage*
Specifies the AMQ message to display.

-M -f *first* -l *last*
Specifies that AMQ messages in a range are displayed from the *first* value to the *last* value.

-r *returnCode*
Specifies the return code to display

-R Specifies that all return codes are displayed.

-R -f *first* -l *last*
Specifies that return codes in a range are displayed from the *first* value to the *last* value.

-s *symbol*
Specifies the symbol to display

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Numeric arguments are interpreted as decimal if they start with a digit 1 - 9, or hex if prefixed with 0x.

- If there is a problem with a message within a range, an indication is displayed before the message text. ? is displayed if there are no matching return codes for the message. ! is displayed if the message severity is not the same as the return code severity.
- For more information about this command in IBM MQ, see mqrc in the IBM MQ documentation.

Examples

- This command displays AMQ message 5005:

  ```
  mqrc AMQ5005
  ```

- This command displays return codes in the range 2505 - 2530:

  ```
  mqrc -R -f 2505 -l 2530
  ```

rcrmqobj

Create a client channel definition table (CCDT) file and place it in a downloadable location.

Syntax

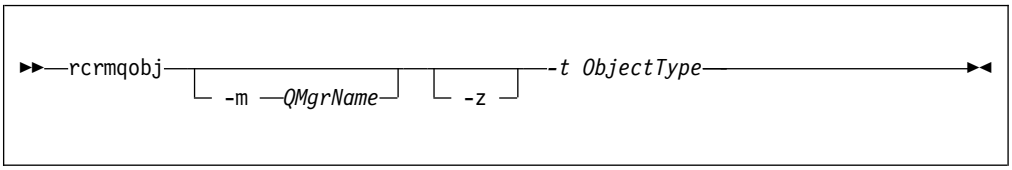

Required parameters

-t ObjectType
The types of object to re-create. The object type for this command when used on the appliance is always `clchltab`.

Optional parameters

-m *QMgrName*
The name of the queue manager for which to re-create objects. If omitted, the command operates on the default queue manager.

-z Suppresses error messages.

rsvmqtrn

Resolve in-doubt and heuristically completed transactions

Purpose

You can use the **rsvmqtrn** command to resolve two different transaction states. You can resolve internal or external in-doubt transactions, and external heuristically completed transactions. Heuristically completed transactions are transactions that are manually resolved by a resource manager but are unacknowledged by the transaction manager.

Syntax

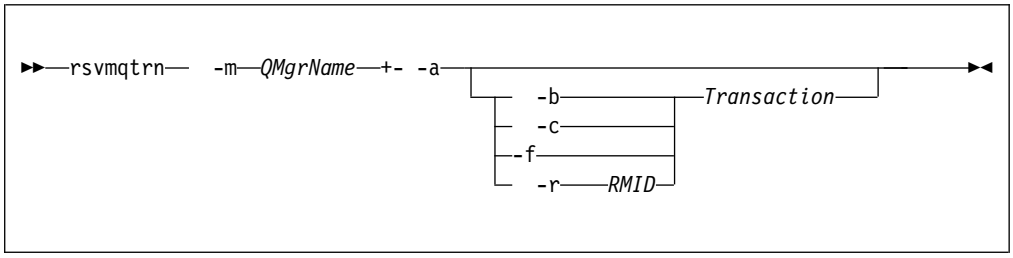

Parameters

-m *QMgrName*
Specifies the name of the queue manager to resolve transactions for.

This parameter is required.

-a Specifies that the queue manager resolves all internally coordinated, in-doubt transactions. That is, all global units of work.

-b *Transaction*
Specifies that the named transaction is backed out.

Transaction specifies the number of the transaction to back out.

This flag is valid for externally coordinated transactions only. That is, for external units of work only.

-c *Transaction*
Specifies that the transaction is committed.

Transaction specifies the number of the transaction to commit.

This flag is valid for externally coordinated transactions only. That is, external units of work only.

-f
Specifies that the named heuristically completed transaction is forgotten.

This flag is valid only for externally coordinated transactions that are resolved, but unacknowledged by the transaction coordinator. That is, external units of work that are resolved, but unacknowledged by the transaction coordinator.

Use this parameter only if the external transaction coordinator is never going to be able to acknowledge the heuristically completed transaction. For example, if the transaction coordinator was deleted.

-r *RMID Transaction*
Specifies that the participation of the resource manager in the in-doubt transaction can be ignored.

The queue manager does not call the resource manager. Instead, it marks the participation of the resource manager in the transaction as being complete.

RMID specifies the ID of the resource manager. *Transaction* specifies the number of the transaction.

This flag is valid for internally coordinated transactions only, and for resource managers that had their resource manager configuration entries removed from the queue manager configuration information.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Use this command only when you are certain that transactions cannot be resolved by the normal protocols. Using this command might result in the loss of transactional integrity between resource managers for a distributed transaction.
- You can use the **dspmqtrn** command to find the number of a transaction.
- For more information about this command in IBM MQ, see rsvmqtrn in the IBM MQ documentation.

Examples

- The following command shows all internally coordinated, in-doubt transactions being resolved for queue manager `QM1`:

```
rsvmqtrn -m QM1 -a
```

Related commands

- "dspmqtrn" on page 341

runmqras

Gather diagnostic information together into a single archive to submit to IBM Support.

Purpose

You can use the **runmqras** command to gather diagnostic information from the appliance into a single archive. You can use this command to gather information about an application or appliance failure, possibly for submission into IBM when you report a problem.

By default, the command gathers information such as the FDC files, error logs, product version, and status information. The command does not gather user information that is contained in messages on queues when you use the default setting. However, if you request sections other than default, the data collected might contain user information.

The diagnostic information is written to a zip file named `runmqras_`*`timestamp`*`.zip`, where *timestamp* has the format *yymmdd_HHMMSS*.

The zip file written to the appliance URI mqdiag://, You can retrieve it by using the copy command (see "**copy**" on page 505), or by using the IBM MQ Appliance web UI (see "Managing files by using the IBM MQ Appliance web UI" on page 233). You can also use the **ftp custom** option of the **runmqras** command to copy the trace directly to an FTP server.

Syntax

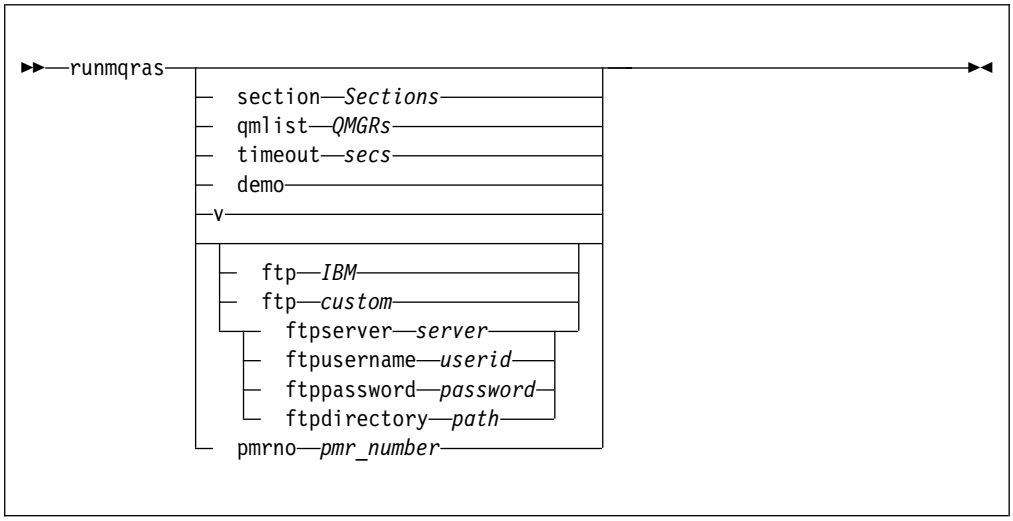

Parameters

section *Sections*

Specifies the optional sections about which to gather more specific information.

By default, a generic section of documentation is collected. Running without requesting more sections is intended as a starting point for general problem diagnosis, but more specific information can be gathered for a specific problem type. You can specify multiple values for *Sections* in a comma-separated list.

IBM support generally supplies you with the sections to use. Example values for *Sections* are the following values:

all Gathers all possible information, including all trace files, and diagnostics for many different types of problems.

This option results in the generation of a very large file, so you must check that the `mqdiag://` directory does not currently contain trace information. If `mqdiag://` does already contain information, you should copy the files off of the appliance, or send them to IBM support, before running **runmqras** with the **all** section.

cluster
Gather information specific for clustering

defs Gather the queue manager definitions and status

nodefault
Prevents the default collections from occurring, but other explicitly requested sections are still collected.

trace Gather all the trace file information plus the default information.

This option results in the generation of a very large file, so you must check that the `mqdiag://` directory does not currently contain trace information. If `mqdiag://` does already contain information, you should copy the files off of the appliance, or send them to IBM support, before running **runmqras** with the **trace** section.

webui A diagnostics test is run on the IBM MQ Console, and the results written to the archive.

qmlist *QMRs*
Specifies one or more queue manager names on which the command is to be run.

You can specify multiple queue managers in a comma-separated list.

By default, the command is run on all queue managers.

timeout *secs*
Specifies the default timeout to give an individual command before the command stops waiting for completion.

A value of zero means that the command waits indefinitely.

The default value is `10`.

demo
Specifies that the command is run in demonstration mode.

In demonstration mode, no commands are processed, and no files gathered. However, you can see which commands would be processed, and which files would be gathered in the `console.log` file that is generated as part of the output.

-v Specifies verbose output.

ftp IBM pmrno *number*
Specifies that the collected archive is sent through basic FTP to IBM.

Anonymous FTP is used to deliver the archive into the IBM ECuRep server. This process is identical to submitting the file manually by using FTP.

number must specify a valid IBM PMR (problem record number) against which to associate the archive.

ftp custom
Specifies that the collected archive is sent through basic FTP to a site of your choosing.

When you use this parameter, you must specify the following *ftp parameters*:

ftpserver *server*
Specifies an FTP server name to connect to.

ftpusername *userid*
Specifies the user ID to log in to the FTP server with.

ftppassword *password*
Specifies the password to log in to the FTP server with

ftpdirectory *path*
Specifies the directory on the FTP server to place the resulting file into.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- For more information about this command in IBM MQ, see runmqras in the IBM MQ documentation.

Examples

- The following command gathers the default documentation from the installation, and all queue managers on the system:

  ```
  runmqras
  ```

- The following command gathers the default documentation from the installation, and sends it directly into IBM to be associated with PMR number 11111,222,333 using the basic FTP capability:

  ```
  runmqras -ftp ibm -pmrno 11111,222,333
  ```

- The following command gathers the default documentation from a machine, plus all trace files, the queue manager definitions, and status for all queue managers on the system:

  ```
  runmqras -section trace,defs
  ```

- The following command copies the information gathered by **runmqras** from the `mqdiag://` directory on the appliance to another location on a system with the IP address 10.10.1.159:

  ```
  (config)# copy mqdiag://runmqras_160818_221406.zip scp://jrb@10.10.1.159//home/user
  ```

runmqsc

Run MQSC commands on a queue manager.

Purpose

You can use the **runmqsc** command to start the `runmqsc` prompt for a queue manager. From the `runmqsc` prompt you can directly enter MQSC commands to perform administration tasks. For example, you can define, alter, or delete a local queue.

Syntax

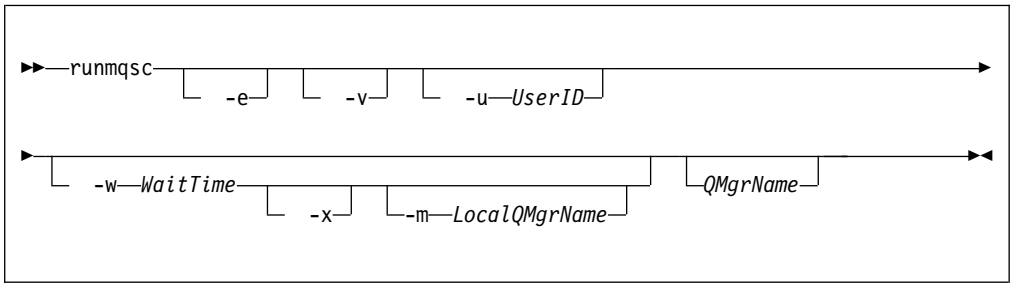

Parameters

-e Specifies that the source text of the MQSC commands is not copied into a report.

This parameter is useful when you enter commands interactively.

-v Specifies that the commands entered are to be verified without performing the action.

You cannot use this parameter with a remote queue manager. That is, the **-w** and **-x** parameters are ignored if specified at the same time as **-v**.

-u *UserID*
Specifies the user ID that the queue manager is accessed with. You are prompted for a matching password.

If you have configured the CONNAUTH AUTHINFO record with CHCKLOCL(REQUIRED) or CHCKLOCL(REQDADM), you must use the **-u** parameter to be able to administer your queue manager with **runmqsc**.

-w ***WaitTime***
Specifies that the MQSC commands are to be run on a remote queue manager. The *WaitTime* specifies how many seconds the command waits for replies from the queue manager. Any replies received after this time are discarded, but the MQSC commands still run.

The *WaitTime* must be a value in the range 1 - 999999.

The replies are received on queue SYSTEM.MQSC.REPLY.QUEUE and the outcome is added to the report. This can be defined as either a local queue or a model queue.

You must have the required channel and transmission queues set up for this. See Preparing channels and transmission queues for remote administration in the IBM MQ documentation.
This parameter is ignored if the **-v** parameter is specified.

-x Specifies that the remote queue manager is running under z/OS. The MQSC commands are then written in a form suitable for the z/OS command queue.

This parameter applies only if the **-w** parameter is also specified.

-m ***LocalQMgrName***
Specifies the local queue manager that you want to use to submit commands to the remote queue manager.

The default value is the local default queue manager.

This parameter applies only if the **-w** parameter is also specified.

QMgrName
Specifies the name of the target queue manager on which to run the MQSC commands.

The default value is the default queue manager.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- To stop the **runmqsc** command, use the **end** command. You can also use the **exit** or the **quit** command.
- For a full list of MQSC commands and their syntax, see The MQSC commands in the IBM MQ documentation.
- The **runmqsc** command takes its input from `stdin`. You can enter MQSC commands interactively by taking `stdin` from the keyboard. Alternatively, you can enter MQSC commands in a file, and run a sequence of frequently used commands by redirecting the input from the file.
- When the commands are processed, the results and a summary are put into a report that is sent to `stdout`. Therefore, you can redirect the output report to a file.

Examples

- The following command starts the `runmqsc` prompt for the default queue manager:

  ```
  runmqsc
  ```

- The following command starts the `runmqsc` prompt for the queue manager `QM1`:

  ```
  runmqsc QM1
  ```

From the `runmqsc` prompt you can directly enter MQSC commands.

runswchl

Switch or query the cluster transmission queues associated with cluster-sender channels.

Purpose

You can use the **runswchl** to switch or query cluster transmission queues that are associated with cluster-sender channels.

The command switches all the stopped or inactive cluster-sender channels that match the `-c` parameter, require switching, and can be switched. The command reports back on the channels that are switched, the channels that do not require switching, and the channels it cannot switch because they are not stopped or inactive.

Syntax

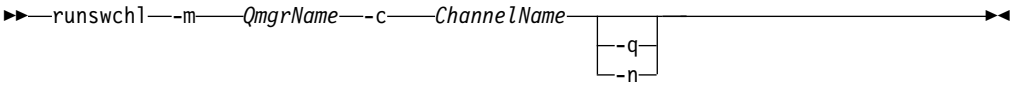

Parameters

-m *QmgrName*
Specifies the queue manager to run the command against.

The queue manager must be started.

This parameter is required.

-c *ChannelName*
Specifies the cluster-sender channels to run the command against.

The *ChannelName* can specify a single channel, or multiple channels if you use a wildcard in the value. You can use an asterisk (*) as a wildcard to match zero or more characters. You can use an asterisk (*) to specify all cluster-sender channels.

This parameter is required.

-q Specifies that the state of the channels is displayed.

If you specify this parameter, no switching occurs. Instead, channels that would be switched are listed.

-n Specifies that when transmission queues are switched, messages are not transferred from the old queue to the new transmission queue.

Messages on the old transmission queue are not transferred unless you associate the transmission queue with another cluster-sender channel.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- For more information about this command in IBM MQ, see runswchl in the IBM MQ documentation.

Examples

- The following command displays the configuration state of cluster-sender channel `TO.QM1`:

  ```
  runswchl -m QM1 -c TO.QM1 -q
  ```

- The following command switches the transmission queue for cluster-sender channel `TO.QM3` without moving the messages on it:

  ```
  runswchl -m QM1 -c TO.QM3 -n
  ```

- The following command switches the transmission queue for cluster-sender channel `TO.QM3` and move the messages on it:

  ```
  runswchl -m QM1 -c TO.QM3
  ```

- The following command displays the configuration state of all cluster-sender channels on `QM1`:

  ```
  runswchl -m QM1 -c * -q
  ```

- The following command displays the configuration state of all cluster-sender channels with a generic name of `TO.*`:

  ```
  runswchl -m QM1 -c TO.* -q
  ```

strmqm

Start a queue manager.

Purpose

You can use the **strmqm** command to start a queue manager.

If the queue manager is part of a high availability configuration, it might start on the other appliance if that is identified as the queue manager's preferred appliance. You can use the status command to check which is the queue manager's preferred appliance. See “status” on page 572.

You can use the **strmqm** command to start a queue manager in the primary role in a disaster recovery configuration. If you use the command to try to start a queue manager in the secondary role, you receive an error message.

Syntax

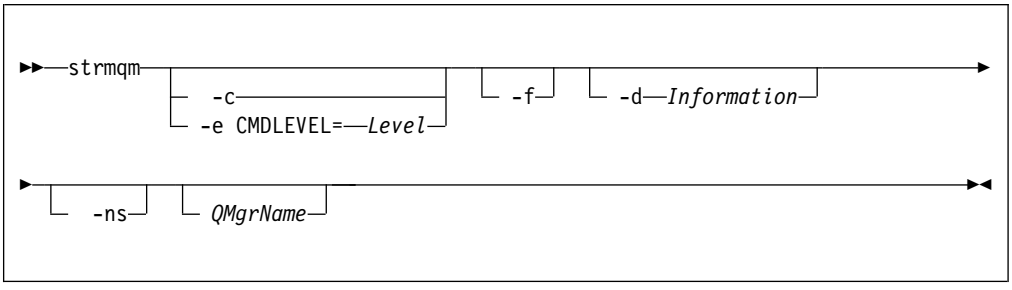

Required parameters

None.

Optional parameters

-c For an HA queue manager, this option only has an effect if it is used with the -ns option.

Specifies that the queue manager default and system objects are to be reset.

Any non-default values for the queue manager default and system objects are replaced with the default values.

The queue manager is stopped after the default and system objects are reset. After you have reset the default and system objects for the queue manager, you must use the **strmqm** command again to start the queue manager.

If you run `mq strmqm -c` on a queue manager that is being used as a IBM MQ Managed File Transfer coordination queue manager, you must rerun the MQSC script that defines the coordination queue manager objects. This script is in a file called *queue_manager_name*`.mqsc`, which is in the IBM MQ Managed File Transfer configuration directory.

-d *Information*
For an HA queue manager, this option only has an effect if it is used with the -ns option.

Specifies whether information messages are displayed.

You can specify one of the following values for *Information*:

all All information messages are displayed.

minimal
The minimal number of information messages are displayed

none No information messages are displayed.

The default value is `all`.

-e CMDLEVEL=*Level*
For an HA queue manager, this option only has an effect if it is used with the -ns option.

Specifies which command level is enabled for the queue manager.

The queue manager is stopped after the command level is set. After you set the command level for the queue manager, you must use the **mq strmqm** command again to start the queue manager.

You must specify a command level that is greater than the current command level of the queue manager and less than or equal to the maximum command level supported by the IBM MQ Appliance.

This flag cannot be specified with `-c`.

-f For an HA queue manager, this option only has an effect if it is used with the -ns option.

Specifies that the queue manager data directory is to be re-created and file permissions are to be reset.

Use this option if you know that a queue manager is not starting because its data directories are missing or corrupted.

If the command is successful, the queue manager starts. If the queue manager fails to start because the configuration information is missing, re-create the configuration information, and restart the queue manager.

You must not use this parameter to re-create the queue manager data directories if you can restore the directories by correcting the configuration. However, you must use this parameter to re-create the queue manager data directory if you are performing media recovery for a queue manager.

-ns
Specifies that the channel initiator, the command server, the listeners, and the services are not started automatically when the queue manager starts. Also specifies that a high availability system does not start. If you start a queue manager that is normally part of an HA configuration, the queue manager starts on the issuing appliance, even if it is not the preferred location. If that appliance subsequently fails, the queue manager will not fail over to the other appliance in the HA group.

QMgrName
Specifies the name of the queue manager to start.

The default value is the default queue manager.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- The **strmqm** command is not required to start a newly-created queue manager. An HA queue manager is started automatically on creation.
- The **strmqm** command is required to restart a stopped HA queue manager. In this case, the queue manager is started on the appliance that is the preferred location for the queue manager, regardless of which appliance the command is issued on. If the HA preferred location is not set, the queue manager starts on the same appliance that it stopped on.
- For more information about creating and activating a backup queue manager, see Backing up and restoring WebSphere MQ queue manager data in the IBM MQ documentation.
- For more information about this command in IBM MQ, see strmqm in the IBM MQ documentation.

Examples

- The following command starts the queue manager `QM1`:

```
strmqm QM1
```

Related commands

- crtmqm (Create queue manager)
- endmqm (Start queue manager)
- dltmqm (Delete queue manager)

strmqtrc

Start trace at a specified level of detail, or report the level of tracing in effect.

Purpose

You can use the **strmqtrc** command to enable tracing. You can specify the tracing that you want:

- You can trace one or more queue managers.
- You can trace one or more processes. The processes can be either part of the product or customer applications that use the IBM MQ API.
- You can trace specific threads within customer applications, either by thread number or by operating system thread number.
- You can trace events. These events can be either the entry or exit from internal functions or the occurrence of a first failure data capture (FDC).
- You can choose from different levels of trace detail.

Trace files are written to the appliance URI mqtrace://, You can retrieve them by using the copy command (see "**copy**" on page 505), or by using the IBM MQ Appliance web UI (see "Managing files by using the IBM MQ Appliance web UI" on page 233)

Syntax

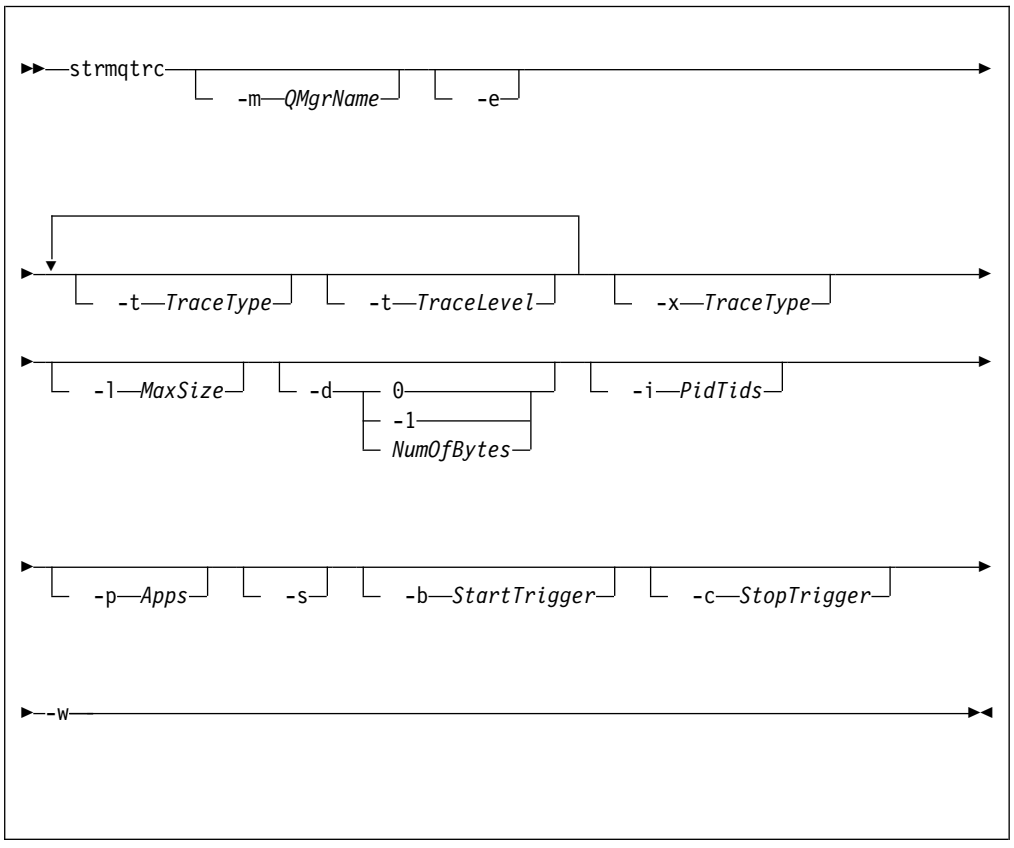

Parameters

-m *QMgrName*

Specifies the name of the queue manager to trace.

You can use an asterisk (*) as a wildcard to replace zero or more characters. You can use a question mark (?) as a wildcard to replace any single character.

-e Specifies that any process that belongs to any component of any queue manager traces its early processing.

You can use this flag to trace the creation or startup of a queue manager.

You cannot use the `-e` flag with the `-m` flag, `-i` flag, the `-p` flag, the `-c` flag, or the `-b` flag.

The default is not to perform early tracing.

-t *TraceType* -t *TraceLevel*
Specifies the points to trace and the amount of trace detail to record.

To specify multiple points to trace, specify multiple `-t` *TraceType* `-t` *TraceLevel* parameters in sequence.

Each *TraceType* can be one of the following values for the points to trace:

all Output data for every trace point in the system. This parameter activates tracing at default detail level.

api Output data for trace points that are associated with the MQI and major queue manager components.

commentary
Output data for trace points that are associated with comments in the components.

comms
Output data for trace points that are associated with data flowing over communications networks.

csdata Output data for trace points that are associated with internal data buffers in common services.

csflows
Output data for trace points that are associated with processing flow in common services.

Explorer
Output data for trace points associated with the IBM MQ Explorer.

Java Output data for trace points associated with applications using the IBM MQ classes for Java™ API.

lqmdata
Output data for trace points that are associated with internal data buffers in the local queue manager.

lqmflows
Output data for trace points that are associated with processing flow in the local queue manager.

otherdata
Output data for trace points that are associated with internal data buffers in other components.

otherflows
Output data for trace points that are associated with processing flow in other components.

parms Activate tracing at default-detail level for flow processing trace points.

remotedata
Output data for trace points that are associated with internal data buffers in the communications component

remoteflows
Output data for trace points that are associated with processing flow in the communications component.

servicedata
Output data for trace points that are associated with internal data buffers in the service component.

serviceflows
Output data for trace points that are associated with processing flow in the service component.

soap Output data for trace points associated with IBM MQ Transport for SOAP.

spldata
Output data for trace points that are associated with buffers and control blocks that use a security policy (AMS) operation.

splflows
Output data for trace points that are associated with entry and exit data for functions that use a security policy (AMS) operation.

ssl Output data that is associated with using GSKit to enable Secure Sockets Layer (SSL) channel security.

versiondata
Output data for trace points that are associated with the version that is running.

The default value is `all`.

Each *TraceLevel* can be one of the following values:

detail Activate tracing at high-detail level for flow processing trace points.

parms Activate tracing at default-detail level for flow processing trace points.

The default value is `parms`.

-x *TraceType*
Specifies the points to exclude from trace.

You can specify the same values for *TraceType* as listed for the `-t` parameter. The default value is `all`.

To specify multiple points to exclude from trace, specify multiple `-x` *TraceType* parameters in sequence.

-l *MaxSize*
Specifies the maximum size of a trace file in megabytes (MB).

The maximum value for *MaxSize* is 2048.

-d 0
Specifies that no user data is traced.

-d -1
Specifies that all user data is traced.

-d *NumOfBytes*
Specifies the number of bytes of data to trace.

For a communication trace, trace the specified number of bytes of data, including the transmission segment header (TSH).

For an MQPUT or MQGET call, trace the specified number of bytes of message data that is held in the message buffer.

Values in the range 1 - 15 are not allowed.

-i *PidTids*
Specifies the process identifier (PID) and thread identifier (TID) to which the trace generation is restricted.

You cannot use the `-i` flag with the `-e` flag.

This parameter must be only used under the guidance of IBM Service personnel.

-p *Apps*
Specifies the named processes to which the trace generation is restricted.

Specify *Apps* as a comma-separated list, with each name in the list specified exactly as the program name would be displayed in the "Program Name" FDC header. You can use an asterisk (*) as a wildcard to match zero or more characters. You can use a question mark (?) to match a single character.

You cannot use the `-p` flag with the `-e` flag.

-s Specifies that the tracing options that are currently in effect are reported.

You must use this parameter on its own with no other parameters.

-b *Start_Trigger*
Specifies the FDC probe IDs for which tracing must be turned on.

Start_Trigger takes one of the following values:

FDC=*comma-separated list of FDC probe IDs*
Turns tracing on when any FDCs with the specified FDC probe IDs are generated.

You can use an asterisk (*) as a wildcard to match zero or more characters. You can use a question mark (?) to match a single character.

You cannot use the `-b` flag with the `-e` flag.

This parameter must be used only under the guidance of IBM Service personnel.

-c *Stop_Trigger*
Specifies the FDC probe IDs for which tracing must be turned off, or interval in seconds after which tracing must be turned off.

Stop_Trigger takes one of the following values:

FDC=*comma-separated list of FDC probe IDs*
Turns tracing off when any FDCs with the specified FDC probe IDs are generated.

You can use an asterisk (*) as a wildcard to match zero or more characters. You can use a question mark (?) to match a single character.

interval=*n*
Where n is an unsigned integer in the range 1 - 32,000,000.

Turns tracing off n seconds after it starts or, if it tracing is already enabled, turns tracing off n seconds after this instance of the command is entered.

This parameter must be used only under the guidance of IBM Service personnel.

-w Allow any application to trigger trace.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Each combination of parameters on an individual invocation of the command are interpreted as having a logical AND between them. You can start the command multiple times, regardless of whether tracing is already enabled. If tracing is already enabled, the trace options that are in effect are modified to those options specified on the most recent invocation of the command.
- Multiple invocations of the command, without an intervening **enqmqtrc** command, are interpreted as having a logical OR between them. The maximum number of concurrent **strmqtrc** commands that can be in effect at one time is 16.
- When a trace file reaches the specified maximum, it is renamed to AMQ*ppppp*.*qq*.TRS and a new AMQ*ppppp*.*qq*.TRC file is started. If a previous copy of an AMQ*ppppp*.*qq*.TRS file exists, it is deleted.
- For more information about this command in IBM MQ, see strmqtrc in the IBM MQ documentation.

Examples

- The following command enables tracing of processing flow from common services and the local queue manager for a queue manager called `exampleQM`. Trace data is generated at the default level of detail.

  ```
  strmqtrc -m exampleQM -t csflows -t lqmflows -t parms
  ```

- The following command disables tracing of SSL activity on a queue manager called `exampleQM`. Other trace data is generated at the `parms` level of detail.

  ```
  strmqtrc -m exampleQM -x ssl -t parms
  ```

- The following command enables high-detail tracing of the processing flow for all components:

  ```
  strmqtrc -t all -t detail
  ```

- The following command enables tracing when FDC KN34650 occurs, and stops tracing when FDC KN346080 occurs. In both cases the FDC must occur on a process that is using queue manager `exampleQM`:

  ```
  strmqtrc -m exampleQM -b FDC=KN346050 -c FDC=KN346080
  ```

The next examples use the -p and -m flags to show how a combination of parameters on an individual invocation of the command are interpreted as having a logical AND between them. The examples also show how multiple invocations of the command, without an intervening **mq enqmqtrc** command, are interpreted as having a logical OR between them:

1. The following command enables tracing for all threads that result from any executing process that is called amqxxx.exe:

   ```
   strmqtrc -p amqxxx.exe
   ```

2.
   - If you start the following command after the command in step 1, without an intervening **endmqtrc** command, then tracing is limited to all threads that result from any running process that is called amqxxx.exe *and* that are using queue manager `exampleQM2`:

     ```
     strmqtrc -p amqxxx.exe -m exampleQM2
     ```

- If you start the following command after the command in step 1, without an intervening **endmqtrc** command, then tracing is limited to all processes and threads that result from running amqxxx.exe *or* that are using queue manager `exampleQM2`:

  ```
  strmqtrc -m exampleQM2
  ```

Related commands

- "endmqtrc" on page 346

IBM MQ configuration commands

You can use the IBM MQ configuration commands to set and clear environments variables for your IBM MQ system.

The queue manager certificate commands can be run from the command line interface in MQ command mode. To enter MQ command mode, type `mqcli`.

dspmqini

Display attributes from the `qm.ini` or `mqat.ini` file of a specified queue manager.

Purpose

You can use the **dspmqini** command to display the `qm.ini` or `mqat.ini` file for a queue manager.

Syntax

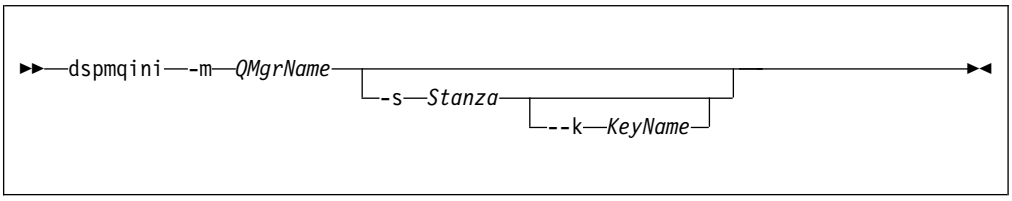

Parameters

-m *QMgrName*
Specifies that the configuration file that is associated with the specified queue manager is displayed.

If you do not specify any further parameters, the command displays all of the stanzas and values for the queue manager.

-s *Stanza*
Specifies which stanza of the file is displayed.

Valid values for *Stanza* are the following values for the `qm.ini` file:

- Log
- TCP
- Channels
- TuningParameters
- SSL
- Security
- Subpool

Valid values for *Stanza* are the following values for the `mqat.ini` file:

- AllActivityTrace

If you do not specify any further parameters, the command displays the values for the stanza specified.

-k ***KeyName***
Specifies which key name of the file is displayed.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command displays the key names and values in the `Log` stanza of the `qm.ini` file of queue manager `QM1`:

```
dspmqini -m QM1 -s Log
```

- The following command displays the value of the key name `ClusterQueueAccessControl` in the `Security` stanza of the `qm.ini` file of queue manager `QM1`:

```
dspmqini -m QM1 -s Security -k ClusterQueueAccessControl
```

Related commands

- setmqini

dspmqvar

Display environment variables set for a specified queue manager.

Purpose

You can use the **dspmqvar** command to display the environment variables that are set for a specified queue manager.

Syntax

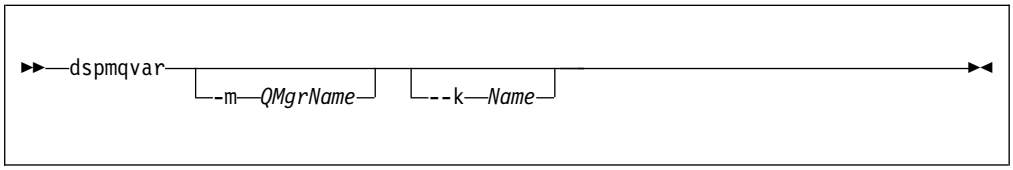

Parameters

-m ***QMgrName***
Specifies the queue manager to display the environment variables for.

If you do not specify this parameter, the global environment variables are displayed.

If you do not specify any further parameters, the command displays all of the environment variables for the queue manager.

-k ***Name***
Specifies which environment variable is displayed.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command displays the value of the environment variable `AMQ_SERVICE_DEBUG_REPOS` for the queue manager `QM1`:

  ```
  dspmqvar -m QM1 -k AMQ_SERVICE_DEBUG_REPOS
  ```

- The following command displays the values of all the global environment variables:

  ```
  dspmqvar
  ```

Related commands

- setmqvar

setmqini

Add or remove an attribute from the `qm.ini` file of a specified queue manager. Set a value for an attribute in the `mqat.ini` file.

Purpose

You can use the **setmqini** command to configure a queue manager by editing the `qm.ini` file for that queue manager. You can also use the command to configure trace levels in the `mqat.ini` file. Changes to the values in the `mqat.ini` file do not take effect until the next time the queue manager is started.

Syntax

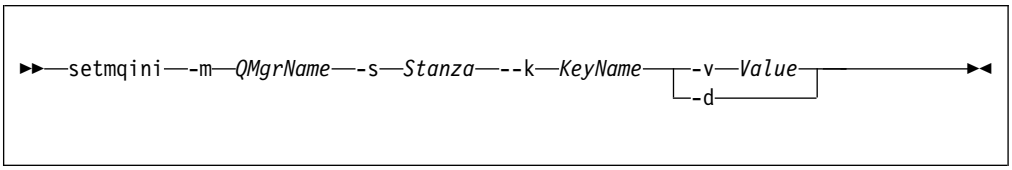

Parameters

-m *QMgrName*
Specifies that the `qm.ini` file that is associated with the specified queue manager is to be modified.

-s *Stanza*
Specifies which stanza of the configuration file is to be added to, or deleted from. The stanza that is specified determines whether the `qm.ini` or `mqat.ini` file is changed.

The following values for *Stanza* modify the `qm.ini` file:

- Log
- TCP
- Channels
- TuningParameters
- SSL

- Security
- Subpool

The following value for *Stanza* modifies the `mqat.ini` file.

- AllActivityTrace

-k *KeyName*
Specifies the key name to add or remove from the `qm.ini` file, or the key to edit in the `mqat.ini` file.

Ensure that the value of *KeyName* is correct before you use the command to add or remove a key and value from the stanza. The value of *KeyName* is not validated. If incorrect values are specified in the `qm.ini` file, a subsequent attempt to start the queue manager might fail.

-v *Value*
Specifies the value to add for the specified key name.

If *Value* is a string that contains spaces, it must be enclosed in double quotation marks. Any double quotation marks that are used in the *Value* must be escaped by using a backslash (\ ).

Ensure that the value of *Value* is correct before you use the command to add or remove a key and value from the stanza. The value of *Value* is not validated. If incorrect values are specified in the `qm.ini` file, a subsequent attempt to start the queue manager might fail.

-d Specifies that the key name specified by the **-k** parameter is deleted from the `qm.ini` file.

To delete an entire stanza, each key name must be deleted individually.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Do not edit the `qm.ini` file to control the number of channels. Instead, use the `MAXINST` and `MAXINSTC` values on your `SVRCONN` channels. For more information, see "Queue manager configuration on the IBM MQ Appliance" on page 19.
- For more information about changing the `qm.ini` file, see Changing queue manager configuration information in the IBM MQ documentation.
- For more information about the `mqat.ini` file, see Configuring activity trace behavior using mqat.ini in the IBM MQ documentation.

Examples

- The following command adds the value `Xmitq` to the key name `RemoteQueueAccessControl` in the `Security` stanza of the `qm.ini` file of queue manager `QM1`:

```
setmqini -m QM1 -s Security -k RemoteQueueAccessControl -v Xmitq
```

- The following command deletes the key name and associated value of `RemoteQueueAccessControl` in the `Security` stanza of the `qm.ini` file of queue manager `QM1`:

```
setmqini -m QM1 -s Security -k RemoteQueueAccessControl -d
```

Related commands

- dspmqini

setmqvar

Add or remove an environment variable for the appliance or for a specified queue manager.

Purpose

You can use the **setmqvar** command to configure environment variables for the appliance or for a specified queue manager.

When you set a specific queue manager variable, the changes take effect the next time that the queue manager is started.

When you set a global environment variable, the changes take effect immediately.

Syntax

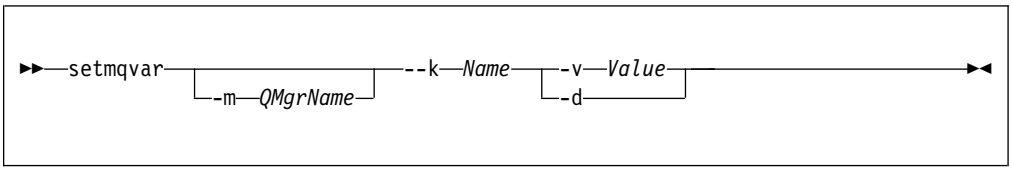

Parameters

-m ***QMgrName***
Specifies the queue manager for which the environment variable is modified.

If this parameter is omitted, the global environment variable is modified.

-k ***Name***
Specifies the name of the environment variable to add or remove.

Ensure that the value of *Name* is correct before you use the command to add or remove an environment variable. The value of *Name* is not validated.

-v ***Value***
Specifies the value to add for the specified environment variable.

If *Value* is a string that contains spaces, it must be enclosed in double quotation marks. Any double quotation marks that are used in the *Value* must be escaped by using a backslash (\ ).

Ensure that the value of *Value* is correct before you use the command to add or remove an environment variable. The value of *Value* is not validated.

-d Specifies that the environment variable specified by the **-k** parameter is deleted.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command adds an environment variable AMQ_SERVICE_DEBUG_REPOS with the value of TRUE to the queue manager QM1:

  ```
  setmqvar -m QM1 -k AMQ_SERVICE_DEBUG_REPOS -v TRUE
  ```

- The following command deletes the global environment variable DEBUG_MODE:

  ```
  setmqvar -k DEBUG_MODE -d
  ```

Related commands

- dspmqvar

Messaging user and group commands

Appliance administrative users use these commands to set up user IDs and groups for messaging users.

These commands must be run from the WebSphere MQ administration mode. If the system is in the WebSphere MQ administration mode the prompt includes mq. To enter the WebSphere MQ administration mode, enter mqcli on the command line. To exit the WebSphere MQ administration mode, enter exit on the command line.

The appliance reserves the following user IDs for its own use:

- hacluster
- mqm
- mqsystem
- root
- sshd

You cannot create user IDs with these names, or delete, modify, or list these user IDs.

The appliance reserves the following groups for its own use:

- haclient
- root
- sshd
- utmp

You cannot create groups with these names, or delete or list these groups.

The appliance also provides the standard IBM MQ mqm group. You cannot delete this group, but you can add users to it.

usercreate

Creates user IDs for messaging users on the IBM MQ Appliance.

Syntax

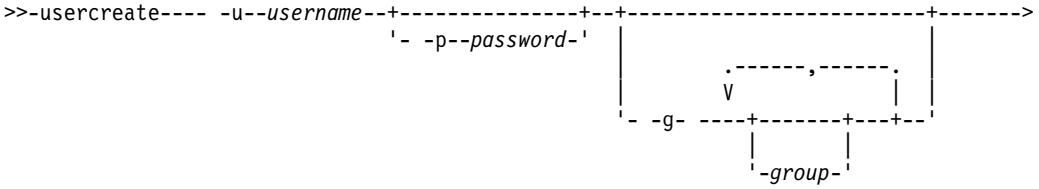

```
>-------+------------------+--->
        '- -d--description-'
```

Parameters

-u *username*
Specifies the user ID to be created.

-p *password*
Optionally specifies a password for the user ID.

-g *group1, group2, groupn...*
Optionally specifies one or more groups that the user belongs to. By default, all new users belong to the group named `users` (you cannot remove users from this group).

-d *text*
Optionally specifies a description for the user ID.

Examples

- The following command creates a new user ID, `myid`, with the password, `pword`, belonging to the admin group:

  ```
  usercreate -u myid -p pword -g MQadmin
  ```

userdelete

Deletes a messaging user on the IBM MQ Appliance.

Syntax

```
►►—userdelete——-u—username————————————————————————————————————►◄
```

Parameters

-u *username*
Specifies the user ID to be deleted.

Examples

- The following command deletes the user ID yourid:

  ```
  userdelete -u yourid
  ```

usermodify

Modifies user IDs for messaging users on the IBM MQ Appliance.

Syntax

```
>>-usermodify---- -u--username--+---------------+--+-------------------------+------->
                                '- -p--password-'  |                         |
                                                   |       .------,------.   |
                                                   |       V             |   |
                                                   '- -g- ----+-------+---+--'
                                                              |       |
                                                              '-group-'


>-------+------------------+--->
        '- -d--description-'
```

Parameters

-u *username*
Specifies the user ID to be modified.

-p *password*
Optionally specifies a password when modifying a user ID.

-g *group1, group2, groupn...*
Optionally specifies one or more groups that the user belongs to when modifying a user ID.

-d *text*
Optionally specifies a description when modifying a user ID.

Examples

- The following command modifies the user, `myid`, changes the password to be `newpword`, and replaces the groups to which the user belongs to `adminUK` and `adminUS`:

  ```
  usermodify -u myid -p newpword -g MQadminUK, MQadminUS
  ```

userlist

Lists the messaging users on the IBM MQ Appliance.

Syntax

Parameters

-u *username*
If you specify a user ID, then the details of that user are listed.

Examples

- The following command lists the current user IDs:

  ```
  userlist
  ```

groupcreate

Adds user groups for messaging users on the IBM MQ Appliance.

Purpose

You can use the **groupcreate** command to work with user groups for the messaging users on the IBM MQ Appliance.

By default, all users belong to the group `users`. You cannot remove users from the `users` group, but you can add them to other groups if required.

Syntax

Parameters

-g *group*
Specifies the user group to be created.

Examples

- The following command creates a new user group, MQgrp:

```
groupcreate -g MQgrp
```

groupdelete

Deletes user groups for messaging users on the IBM MQ Appliance.

Purpose

You can use the **groupdelete** command to work with user groups for the messaging users on the IBM MQ Appliance.

Syntax

Parameters

-g *group*
Specifies the user group to be deleted.

Examples

- The following command deletes the user group, MQ0grp:

```
groupdelete -g MQ0grp
```

grouplist

Lists user groups for messaging users on the IBM MQ Appliance.

Purpose

You can use the **grouplist** command to work with user groups for the messaging users on the IBM MQ Appliance.

Syntax

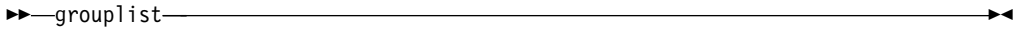

Examples

- The following command lists the user groups on the appliance:

```
grouplist
```

userbackup

Backs up messaging users on the IBM MQ Appliance to a file.

Purpose

You can use the **userbackup** command to create a back up file containing the details for messaging users that have been defined on the IBM MQ Appliance.

Syntax

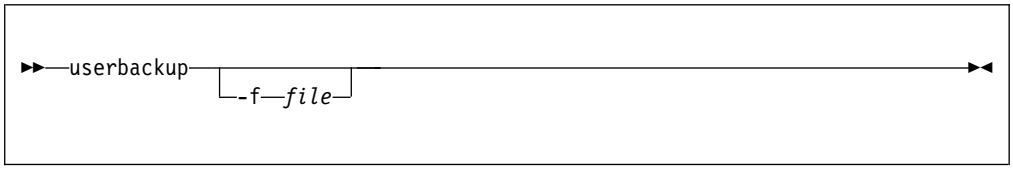

Parameters

-f *filename*
Specifies the file that the messaging users are backed up to. The file is written to the `mqbackup:` location on the IBM MQ Appliance. If you do not specify a file name, the back up is written to a file named `userbackup-`*`date-time`*, for example, `userbackup-20150219-132655`.

Examples

- The following command backs up users to the file `backup_15115`:

```
userbackup -f backup_15115
```

userrestore

Restores messaging users on the IBM MQ Appliance from a file to which they were previously backed up.

Purpose

You can use the **userrestore** command to restore messaging users from back up file containing the details for messaging users.

Syntax

Parameters

-f *filename*
Specifies the file containing the users to restore. The file must be located in the `mqbackup:` location on the IBM MQ Appliance.

Examples

- The following command restores users from the file `backup_15115`:

```
userrestore -f backup_15115
```

Queue manager certificate management commands

You can use the certificate management commands to manage certificates for queue managers.

The queue manager certificate commands can be run from the command line interface in MQ command mode. To enter MQ command mode, type `mqcli`.

addcert

Add the public part of a certificate to the keystore of a specific queue manager.

Purpose

You can use the **addcert** command to add the public part of a certificate to the key repository of a specified queue manager.

Syntax

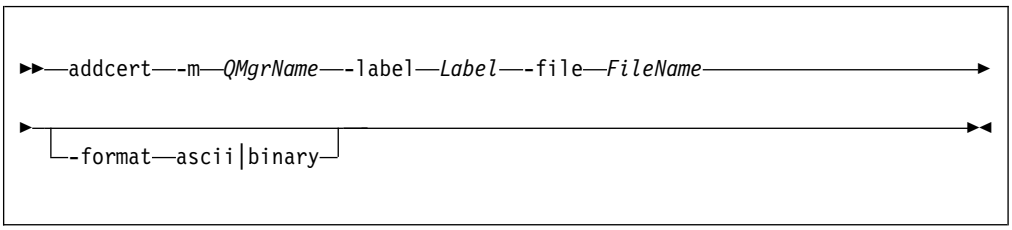

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the certificate is added.

The queue manager must exist.

-label *Label*
Specifies the label that is associated with the certificate.

For more information about valid syntax for the certificate label, see Digital certificate labels, understanding the requirements in the IBM MQ documentation.

-file *FileName*
Specifies the file that contains the certificate.

This file must be available on the appliance. The file must be located in `mqpubcert://`.

-format ascii|binary
Specifies the format of the certificate.

The default value is ascii.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command adds a CA certificate from file `CA.pem`, with a label of `CACert` to the key repository for the queue manager `QM1`:

```
addcert -m QM1 -file CA.pem -label CACert
```

Related commands

- “createcert” on page 376
- “deletecert” on page 379
- “detailcert” on page 381

- "listcert" on page 385
- "receivecert" on page 387

createcert

Create a self-signed certificate for a queue manager.

Purpose

You can use the **createcert** command to create a self-signed certificate and add it to the key repository of a specified queue manager. The certificate data is extracted from the newly created certificate and placed in a file.

Syntax

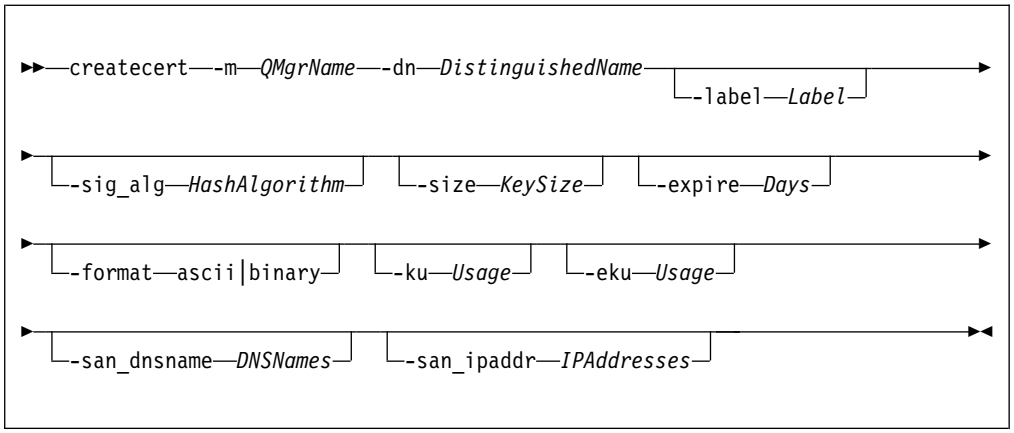

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the self-signed certificate is created.

The queue manager must exist.

-dn *DistinguishedName*
Specifies the X.500 distinguished name that uniquely identifies the certificate.

DistinguishedName is a string that is enclosed in double quotation marks. For example, `"CN=John Smith,O=IBM,OU=Test,C=GB"`. The CN, O, and C attributes are required.

-label *Label*
Specifies the label that is associated with the certificate.

The default value is ibmwebspheremq<*QMgrName*>, where *QMgrName* is the name of the queue manager in lowercase.

-sig_alg *HashAlgorithm*
Specifies the signing algorithm that is used to create the signature that is associated with the new self-signed certificate.

HashAlgorithm can be one of the following values:
`md5`, `MD5_WITH_RSA`, `MD5WithRSA`, `SHA_WITH_DSA`, `SHA_WITH_RSA`, `sha1`, `SHA1WithDSA`, `SHA1WithECDSA`, `SHA1WithRSA`, `sha224`, `SHA224_WITH_RSA`, `SHA224WithDSA`, `SHA224WithECDSA`, `SHA224WithRSA`, `sha256`, `SHA256_WITH_RSA`, `SHA256WithDSA`, `SHA256WithECDSA`, `SHA256WithRSA`, `SHA2WithRSA`, `sha384`, `SHA384_WITH_RSA`, `SHA384WithECDSA`, `SHA384WithRSA`, `sha512`, `SHA512_WITH_RSA`, `SHA512WithECDSA`,

SHA512WithRSA, SHAWithDSA, SHAWithRSA , EC_ecdsa_with_SHA1, EC_ecdsa_with_SHA224, EC_ecdsa_with_SHA256, EC_ecdsa_with_SHA384, or EC_ecdsa_with_SHA512.

The default value is SHA256WithRSA.

-size *KeySize*
Specifies the size of the new key pair.

The default value is 2048 for all RSA signature algorithms. For elliptic curve signature algorithms, use 256, 384, or 512 to match the selected algorithm.

-expire *Days*
Specifies the expiration time of the certificate, in days.

The default value is 365.

-format ascii|binary
Specifies the format of the output file.

The default value is ascii.

-ku *Usage*
Specifies a list of valid uses for the certificate.

To specify more than one use, enter each value in a comma-separated list.

-eku *Usage*
Specifies a list of valid uses for the certificate.

To specify more than one use, enter each value in a comma-separated list.

-san_dnsname *DNSNames*
Specifies the Subject Alternative Name (SAN) DNS names for the certificate that is created.

To specify more than one DNS name, enter each value in a comma-separated list.

-san_ipaddr *IPAddresses*
Specifies the Subject Alternative Name (SAN) IP addresses for the certificate that is created.

To specify more than one IP address, enter each value in a comma-separated list.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mq. To enter the IBM MQ administration mode, enter mqcli on the command line. To exit the IBM MQ administration mode, enter exit on the command line.
- The target file name is generated based on the label that is specified in the command. The file name is displayed when the command completes.

Examples

- The following command creates a certificate for queue manager QM1, with a distinguished name of "CN=John Smith,O=IBM,OU=Test,C=GB":

```
createcert -m QM1 -dn "CN=John Smith,O=IBM,OU=Test,C=GB"
```

Related commands

- “addcert” on page 375
- “deletecert” on page 379

- "detailcert" on page 381
- "listcert" on page 385
- "receivecert" on page 387

createcertrequest

Create a certificate request for a queue manager.

Purpose

You can use the **createcertrequest** command to create a certificate request for a specified queue manager.

Syntax

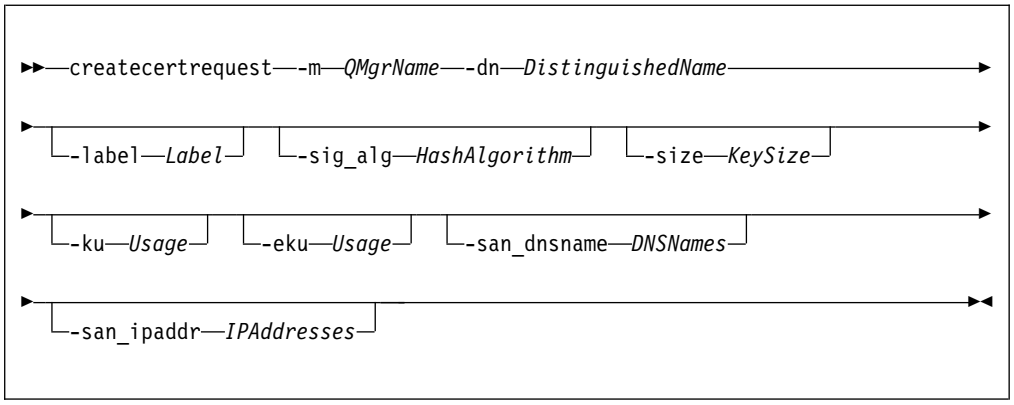

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the certificate request is created.

The queue manager must exist.

-dn *DistinguishedName*
Specifies the X.500 distinguished name that uniquely identifies the certificate.

DistinguishedName is a string that is enclosed in double quotation marks. For example, `"CN=John Smith,O=IBM,OU=Test,C=GB"`. The CN, O, and C attributes are required.

-label *Label*
Specifies the label that is associated with the certificate request.

The default value is ibmwebspheremq<*QMgrName*>, where *QMgrName* is the name of the queue manager in lowercase.

-sig_alg *HashAlgorithm*
Specifies the signing algorithm that is used to create the signature that is associated with the new certificate.

HashAlgorithm can be one of the following values:
`md5`, `MD5_WITH_RSA`, `MD5WithRSA`, `SHA_WITH_DSA`, `SHA_WITH_RSA`, `sha1`, `SHA1WithDSA`, `SHA1WithECDSA`, `SHA1WithRSA`, `sha224`, `SHA224_WITH_RSA`, `SHA224WithDSA`, `SHA224WithECDSA`, `SHA224WithRSA`, `sha256`, `SHA256_WITH_RSA`, `SHA256WithDSA`, `SHA256WithECDSA`, `SHA256WithRSA`, `SHA2WithRSA`, `sha384`, `SHA384_WITH_RSA`, `SHA384WithECDSA`, `SHA384WithRSA`, `sha512`, `SHA512_WITH_RSA`, `SHA512WithECDSA`,

SHA512WithRSA, SHAWithDSA, SHAWithRSA , EC_ecdsa_with_SHA1, EC_ecdsa_with_SHA224, EC_ecdsa_with_SHA256, EC_ecdsa_with_SHA384, or EC_ecdsa_with_SHA512.

The default value is SHA256WithRSA.

-size *KeySize*
Specifies the size of the new key pair.

The default value is 2048 for all RSA signature algorithms. For elliptic curve signature algorithms, use 256, 384, or 512 to match the selected algorithm.

-ku *Usage*
Specifies a list of valid uses for the certificate.

To specify more than one use, enter each value in a comma-separated list.

-eku *Usage*
Specifies a list of valid uses for the certificate.

To specify more than one use, enter each value in a comma-separated list.

-san_dnsname *DNSNames*
Specifies the Subject Alternative Name (SAN) DNS names for the certificate that is created.

To specify more than one DNS name, enter each value in a comma-separated list.

-san_ipaddr *IPAddresses*
Specifies the Subject Alternative Name (SAN) IP addresses for the certificate that is created.

To specify more than one IP address, enter each value in a comma-separated list.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mq. To enter the IBM MQ administration mode, enter mqcli on the command line. To exit the IBM MQ administration mode, enter exit on the command line.
- The certificate request file name is generated based on the label that is specified in the command. The file name is displayed when the command completes.

Examples

- The following command creates a certificate request for queue manager QM2, with a distinguished name of "CN=Jane Smith,O=IBM,OU=Test,C=US":

```
createcertrequest -m QM2 -dn "CN=Jane Smith,O=IBM,OU=Test,C=US"
```

Related commands

- “deletecertrequest” on page 380
- “detailcertrequest” on page 382
- “listcertrequest” on page 386
- “recreatecertrequest” on page 388

deletecert

Delete a certificate from the keystore of a specific queue manager.

Purpose

You can use the **deletecert** command to remove a certificate from the key repository of a specified queue manager.

Syntax

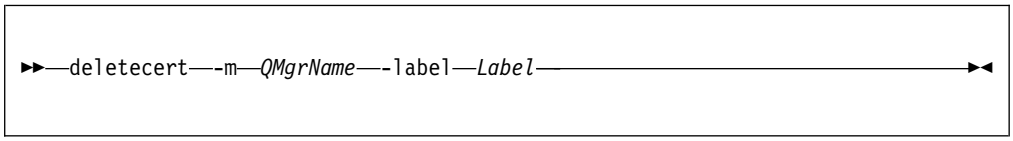

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the certificate is deleted.

The queue manager must exist.

-label *Label*
Specifies the label that is associated with the certificate.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command deletes a certificate file `certificate.pem`, with a label of `cert1` for the queue manager `QM1`:

```
deletecert -m QM1 -file certificate.pem -label cert1
```

Related commands

- "addcert" on page 375
- "createcert" on page 376
- "detailcert" on page 381
- "listcert" on page 385
- "receivecert" on page 387

deletecertrequest

Delete a certificate request that was previously issued from a specific queue manager.

Purpose

You can use the **deletecertrequest** command to delete a certificate request that was previously issued from a specified queue manager.

Syntax

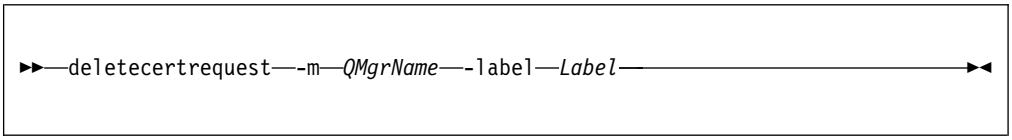

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the certificate request is deleted.

The queue manager must exist.

-label *Label*
Specifies the label that is associated with the certificate request.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command deletes a certificate request with a label of `cert1` for the queue manager `QM1` :

```
deletecert -m QM1 -label cert1
```

Related commands

- “createcertrequest” on page 378
- “detailcertrequest” on page 382
- “listcertrequest” on page 386
- “recreatecertrequest” on page 388

detailcert

Show detailed information about a certificate for a specific queue manager.

Purpose

You can use the **detailcert** command to show detailed information about a specific certificate, or about the queue manager default certificate

Syntax

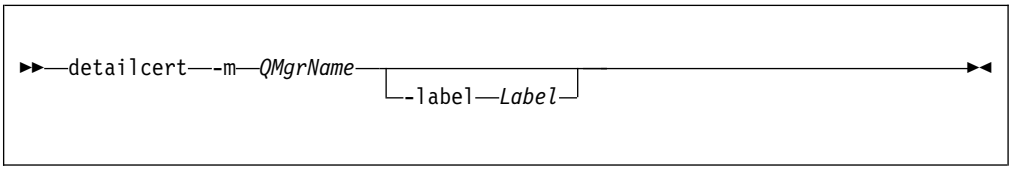

Parameters

-m ***QMgrName***
Specifies the name of the queue manager for which the certificate details are shown.

The queue manager must exist.

-label ***Label***
Specifies the label of the certificate for which detailed information is shown.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command shows the details of the certificate `ibmwebspheremqqm1` for the queue manager `QM1` :

```
detailcert -m QM1 -label ibmwebspheremqqm1
```

Related commands

- "addcert" on page 375
- "createcert" on page 376
- "deletecert" on page 379
- "listcert" on page 385
- "receivecert" on page 387

detailcertrequest

Show detailed information about a certificate request for a specific queue manager.

Purpose

You can use the **detailcertrequest** command to show detailed information about a certificate request for a specified queue manager.

Syntax

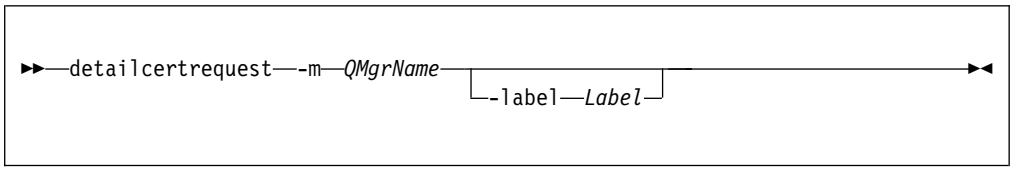

Parameters

-m ***QMgrName***
Specifies the name of the queue manager for which the certificate request details are shown.

The queue manager must exist.

-label ***Label***
Specifies the label of the certificate request for which detailed information is shown.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command shows the details of the certificate request `request1` for the queue manager `QM1` :

```
detailcertrequest -m QM1 -label request1
```

Related commands

- “createcertrequest” on page 378
- “deletecertrequest” on page 380
- “listcertrequest” on page 386
- “recreatecertrequest” on page 388

keybackup

Back up the queue manager key repository to a file.

Purpose

You can use the **keybackup** command to write a copy of the queue manager key repository to a file. You can then copy the file to another system, and restore it when required.

The command creates a compressed archive (.tar.gz) of the key repository files This includes the .kdb and .rdb files, and the crl file, if present. It does not include the password stash file. At completion the name of the archive file, and the password that was stored in the password stash file is displayed. The password is needed to restore the key repository.

Syntax

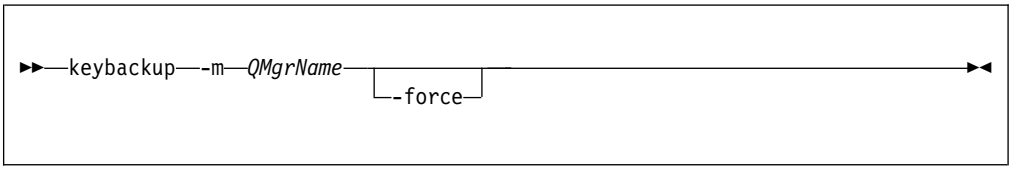

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the key repository is backed up.

The queue manager must exist.

-force
Forces the back up, without displaying a warning about the security issues raised by backing up the key repsotory.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter

the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

- The operation could be regarded as insecure as it places a copy of the queue manager Key Repository into the user accessible file area on the appliance. Unless you specify the -force parameter, the appliance prompts you to confirm that you want to continue with the back up:

```
This operation will generate a copy of your queue manager key  repository,
    which may include private keys.  Although encrypted, you  should take appropriate security
    precautions in handling this file.  The  password required if you ever need to modify or
    restore this file will  be displayed after the copy has been created.  Do you wish to continue
    [Y/N]
```

Examples

- The following command backs up the key repository for the queue manager `QM1`:

```
keybackup -m QM1
```

Related commands

- "keyrestore"
- "**copy**" on page 505

keyrestore

Restore a key repository

You can use the **keyrestore** command to restore to a queue manager a key repository that you have previously backed up.

This command will prompt for password unless one is provided, and then replace the .kdb, .rdb, and .crl (if present) files for this queue manager with the content of the archive file provided. It will then generate a new password stash file.

Syntax

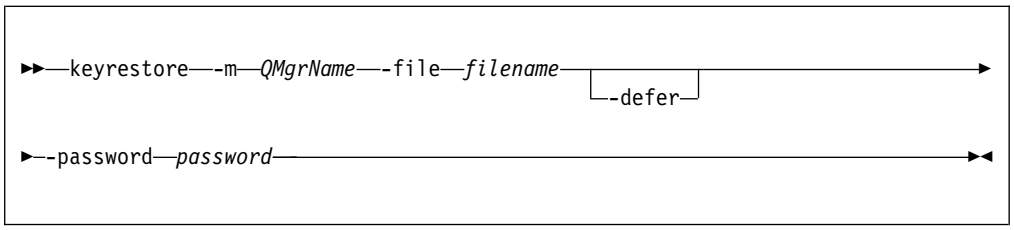

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the key repository is backed up.

The queue manager must exist.

-file *filename*
Specifies the name of the archive file containing the key repository that you are restoring.

-defer
By default, the key repository is restored to the queue manager immediately. If you specify the **-defer** parameter, the action is suppressed until an administrator has manually stopped SSL/TLS channels on that queue manager, and issued a **MQSC REFRESH SECURITY TYPE(SSL)** command.

-password *password*
When running the **keyrestore** command, you must specify the password that was displayed when the archive was created using the **keybackup** command.

You must enclose the password in double quotes if it includes special characters. You must also escape any backslash or double quote characters that are part of the password with a backslash character. For example. if the **keybackup** command returned `pass"word\`, then you should supply the password to the **keyrestore** command as shown:

```
"pass\"word\\"
```

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command restores the key repository for the queue manager `QM1`:

```
keyrestore -m QM1 -file QM1keystore.tar.gz
```

Related commands

- "keybackup" on page 383
- "**copy**" on page 505

listcert

List the certificates that are held in the keystore of a specific queue manager.

Purpose

You can use the **listcert** command to list the certificates that are held in the key repository of a queue manager.

Syntax

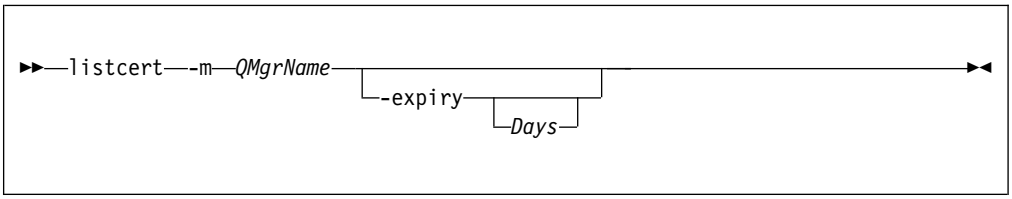

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the certificates are listed.

The queue manager must exist.

-expiry *Days*
Specifies that the valid-from and valid-to dates are displayed.

Days is an optional numeric value that specifies that the valid-from and valid-to dates are displayed for certificates that expire within that number of days.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command lists the certificates for the queue manager `QM1`:

```
listcert -m QM1
```

Related commands

- "addcert" on page 375
- "createcert" on page 376
- "deletecert" on page 379
- "detailcert" on page 381
- "receivecert" on page 387

listcertrequest

List the certificate requests that are outstanding in the keystore of a specific queue manager.

Purpose

You can use the **listcertrequest** command to list the certificate requests that are outstanding in the key repository of a specified queue manager.

Syntax

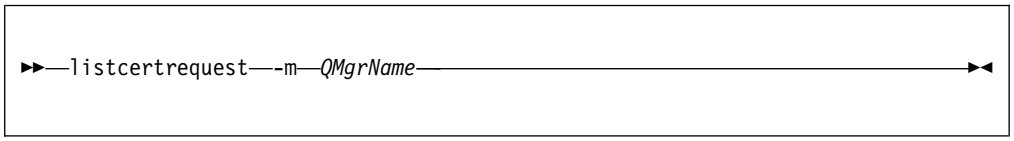

Parameters

-m ***QMgrName***

Specifies the name of the queue manager for which the oustanding certificate requests are listed.

The queue manager must exist.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command lists the outstanding certificate requests for the queue manager `QM1` :

```
listcertrequest -m QM1
```

Related commands

- “createcertrequest” on page 378
- “deletecertrequest” on page 380
- “detailcertrequest” on page 382
- “recreatecertrequest” on page 388

receivecert

Receive a certificate signed by a Certificate Authority (CA) as the result of a previous request.

Purpose

You can use the **receivecert** command to accept a certificate that has been signed by a CA.

Syntax

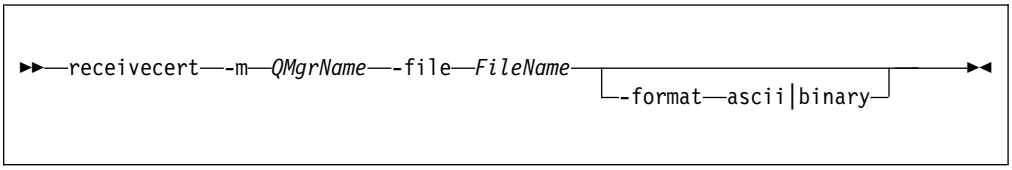

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the certificate is accepted.

The queue manager must exist.

-file *FileName*
Specifies the file that contains the certificate.

This file must be available on the appliance. The file must be located in `mqpubcert://`.

-format ascii|binary
Specifies the format of the certificate.

The default value is ascii.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command accepts a certificate file `certificate.pem` for the queue manager `QM1` :

```
receivecert -m QM1 -file certificate.pem
```

Related commands

- “addcert” on page 375
- “createcert” on page 376
- “deletecert” on page 379

- "detailcert" on page 381
- "listcert" on page 385

recreatecertrequest

Recreate a certificate request for a specific queue manager.

Purpose

You can use the **recreatecertrequest** command to recreate a certificate request. The certificate request can then be sent to a CA to be renewed.

Syntax

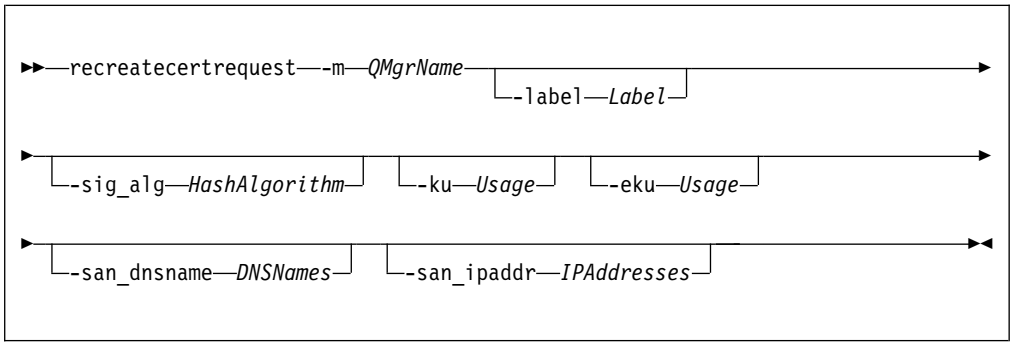

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which the certificate request is re-created.

The queue manager must exist.

-label *Label*
Specifies the label that is associated with the certificate request.

The default value is ibmwebspheremq<*QMgrName*>, where *QMgrName* is the name of the queue manager in lowercase.

-sig_alg *HashAlgorithm*
Specifies the signing algorithm that is used to create the signature that is associated with the new certificate.

HashAlgorithm can be one of the following values:
`md5`, `MD5_WITH_RSA`, `MD5WithRSA`, `SHA_WITH_DSA`, `SHA_WITH_RSA`, `sha1`, `SHA1WithDSA`, `SHA1WithECDSA`, `SHA1WithRSA`, `sha224`, `SHA224_WITH_RSA`, `SHA224WithDSA`, `SHA224WithECDSA`, `SHA224WithRSA`, `sha256`, `SHA256_WITH_RSA`, `SHA256WithDSA`, `SHA256WithECDSA`, `SHA256WithRSA`, `SHA2WithRSA`, `sha384`, `SHA384_WITH_RSA`, `SHA384WithECDSA`, `SHA384WithRSA`, `sha512`, `SHA512_WITH_RSA`, `SHA512WithECDSA`, `SHA512WithRSA`, `SHAWithDSA`, `SHAWithRSA` , `EC_ecdsa_with_SHA1`, `EC_ecdsa_with_SHA224`, `EC_ecdsa_with_SHA256`, `EC_ecdsa_with_SHA384`, or `EC_ecdsa_with_SHA512`.

The default value is `SHA256WithRSA`.

-ku *Usage*
Specifies a list of valid uses for the certificate.

To specify more than one use, enter each value in a comma-separated list.

-eku *Usage*
Specifies a list of valid uses for the certificate.

To specify more than one use, enter each value in a comma-separated list.

-san_dnsname *DNSNames*
Specifies the Subject Alternative Name (SAN) DNS names for the certificate that is created.

To specify more than one DNS name, enter each value in a comma-separated list.

-san_ipaddr *IPAddresses*
Specifies the Subject Alternative Name (SAN) IP addresses for the certificate that is created.

To specify more than one IP address, enter each value in a comma-separated list.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `mq`. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command recreates a certificate request for the queue manager `QM1` :

```
recreatecertrequest -m QM1
```

Related commands

- “createcertrequest” on page 378
- “deletecertrequest” on page 380
- “detailcertrequest” on page 382
- “listcertrequest” on page 386

Queue manager commands

You can use the queue manager commands to create, delete, and manage queue managers.

The queue manager commands can be run from the command line interface in MQ command mode. To enter MQ command mode, type `mqcli`.

See “Differences between queue managers that are running on the IBM MQ Appliance and an IBM MQ installation” on page 18 for specific information about using commands on the appliance.

crtmqm

Create a queue manager.

Purpose

You can use the **crtmqm** command to create a queue manager.

Syntax

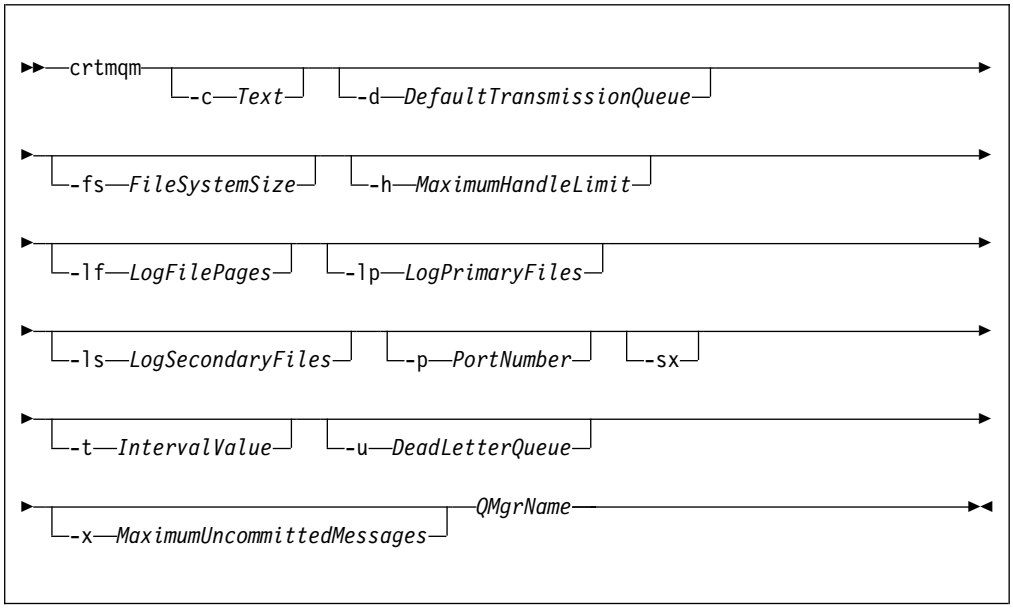

Parameters

QMgrName
Specifies the name of the queue manager that you want to create.

The queue manager name must be the last parameter that is specified in the command.

The name can contain up to 48 characters. The following characters can be used:
`0-9 A-Z a-z . / _ %`

The name of the queue manager must be unique on the IBM MQ Appliance. If the queue manager connects to other queue managers, the queue manager names must be unique within that group of queue managers.

This parameter is required.

-c *Text*
Specifies descriptive text for this queue manager.

You can use up to 64 characters. If you include special characters, enclose the description in single quotation marks. The maximum number of characters is reduced if you are using a double-byte character set (DBCS).

The default value is all blanks.

-d *DefaultTransmissionQueue*
Specifies the name of the local transmission queue where remote messages are put if a transmission queue is not explicitly defined for their destination.

There is no default value.

-fs *FileSystemSize*
Specifies that the queue manager is created with the file system size *FileSystemSize*. If you do not specify this argument, the file system size defaults to 64 GB.

FileSystemSize is a numeric value, which is specified in GB. You can specify a value in MB by entering the value followed by the character `M`. For example, to specify a *FileSystemSize* of 3 GB, enter `3`. To specify a *FileSystemSize* of 1024 MB, enter `1024M`.

For the appliance the minimum value is 128 MB.

Note: After a queue manager is created you cannot resize the file system; ensure the value that is specified here is sufficient for the current and any future workload.

-h ***MaximumHandleLimit***
Specifies the maximum number of handles that an application can open at the same time.

Specify a value in the range 1 - 999999999.

The default value is 256.

-lf *LogFilePages*
Specifies the number of log file pages to use for the log files.

The log data is held in a series of files called log files. The log file size is specified in units of 4 KB pages.

The default number of log file pages is 4096, giving a log file size of 16 MB. The minimum number of log file pages is 64 and the maximum is 65535.

-lp *LogPrimaryFiles*
Specifies the log files that are allocated when the queue manager is created.

The minimum number of primary log files you can have is 2 and the maximum is 510. The default is 3. The total number of primary and secondary log files must not exceed 511 and must not be less than 3.

You can change this value after the queue manager is created. However, the change is not effective until the queue manager is restarted.

-ls *LogSecondaryFiles*
Specifies the log files that are allocated when the primary files are exhausted.

The minimum number of secondary log files you can have is 2 and the maximum is 509. The default is 2. The total number of primary and secondary log files must not exceed 511 and must not be less than 3.

You can change this value after the queue manager is created. However, the change is not effective until the queue manager is restarted.

-p ***PortNumber***
Create a managed TCP listener on the specified port.

Specify a valid port value to create a TCP listener object that uses the specified port. The new listener is called SYSTEM.LISTENER.TCP.1. This listener is under queue manager control, and is started and stopped along with the queue manager.

-sx
Specifies that the queue manager is a high availability (HA) queue manager.

The queue manager starts automatically as part of the HA group.

-t ***IntervalValue***
Specifies the trigger time interval in milliseconds for all queues that are controlled by this queue manager.

That is, after the queue manager receives a trigger-generating message, triggering is suspended for the length of time that is specified by *IntervalValue*.

Specify a value in the range 0 - 999999999.

The default value is 999999999 milliseconds. This value effectively means that triggering is disabled after the first trigger message.

-u *DeadLetterQueue*
Specifies the name of the local queue that is to be used as the dead-letter (undelivered-message) queue.

The default is no dead-letter queue.

-x *MaximumUncommittedMessages*
Specifies the maximum number of uncommitted messages under any one sync point.

The uncommitted messages are the sum of the following messages:

- The number of messages that can be retrieved from queues
- The number of messages that can be put on queues
- Any trigger messages that are generated within this unit of work

The limit that is specified does not apply to messages that are retrieved or put outside a sync point.

Specify a value in the range 1 - 999999999.

The default value is 10000 uncommitted messages.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- After you create the queue manager, you can use the **strmqm** command to start the queue manager. A high availability queue manager is started automatically after creation, so you do not need to start it by using **strmqm**.
- When a queue manager is created, the default and system objects are also created. These objects are listed in System and default objects in the IBM MQ documentation.
- For more information about this command in IBM MQ, see ctrmqm in the IBM MQ documentation.

Examples

- The following command creates a queue manager that is called `QM1`, with a description of `example queue manager`, and creates the system and default objects:

  ```
  crtmqm -c "example queue manager" QM1
  ```

- The following command creates a queue manager that is called `QM2`. It creates the system and default objects, sets the trigger interval to 5000 milliseconds (5 seconds), and specifies `SYSTEM.DEAD.LETTER.QUEUE` as its dead-letter queue.

  ```
  crtmqm -t 5000 -u SYSTEM.DEAD.LETTER.QUEUE QM2
  ```

Related commands

- strmqm (Start queue manager)
- endmqm (End queue manager)
- dltmqm (Delete queue manager)

dltmqm

Delete a queue manager.

Purpose

You can use the **dltmqm** command to delete a queue manager.

Syntax

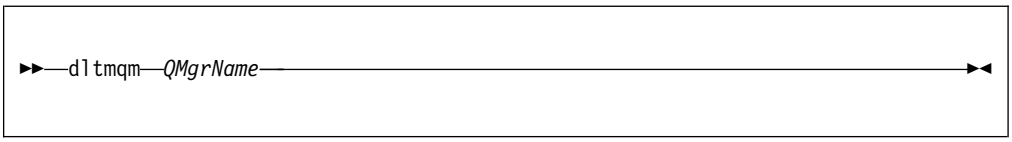

Parameters

QMgrName

Specifies the name of the queue manager that you want to delete.

This parameter is required.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Before you delete the queue manager, you must end the queue manager by using the **endmqm** command.
- For more information about this command in IBM MQ, see dltmqm in the IBM MQ documentation.

Examples

- The following command deletes the queue manager `QM1`.

```
dltmqm QM1
```

Related commands

- crtmqm (Create queue manager)
- strmqm (Start queue manager)
- endmqm (Delete queue manager)

dmpmqcfg

Dump the configuration of a queue manager.

Purpose

You can use the **dmpmqcfg** command to dump the configuration of a queue manager.

Syntax

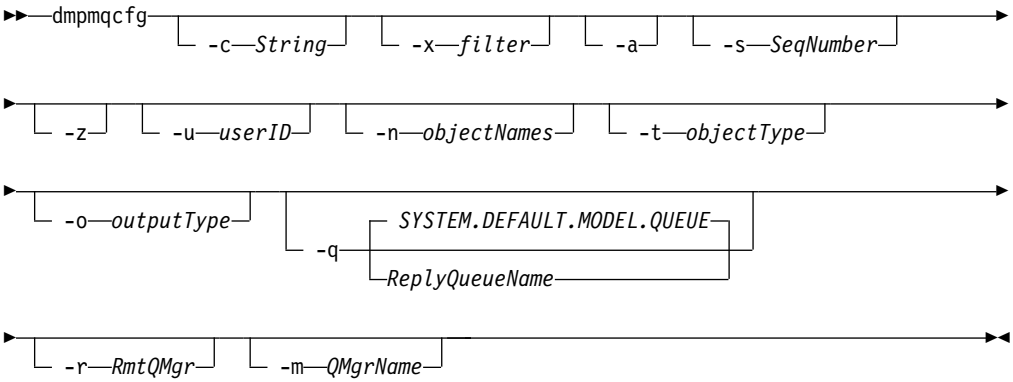

Parameters

-c *string*
Specifies that a client mode connection is used to connect to the queue manager.

string can take one of the following values:

default
Specifies that the default client connection process is used.

"DEFINE CHANNEL(*chlname*) CHLTYPE(CLNTCONN) CONNAME('*connname*')"
Specifies that the specific client channel specified by *chlname* is used to connect to the queue manager at *connname*.

connname specifies the location of the queue manager in the following format *host*(*portnumber*)

If **-c** is omitted, the command connects to the queue manager by using server bindings. If that connection fails, client bindings are used.

-x *filter*
Specifies that the procedure is filtered.

filter can be one of the following values:

object

authority records

channel authentication

subscriptions

all

The default value is `all`.

-a Specifies that object definitions show all attributes.

The default is to return only attributes that differ from the defaults for the object type.

-s *SeqNumber*
Specifies that the channel sequence number for sender, server, and cluster sender channel types is reset to the value specified.

SeqNumber must be in the range 1 - 999999999.

-z Specifies that the command runs in silent mode.

All warnings, such as those that appear when attributes from a queue manager of a higher command level are inquired, are suppressed.

-n *objectNames*
Specifies that the definitions produced by object or profile name are filtered.

The object or profile name can contain a single asterisk. The * option can be placed only at the end of the entered filter string.

-t *objectType*
Specifies a single type of object to export.

objectType can be one of the following values:

all All object types.

authinfo
An authentication information object.

channel
A channel (including MQTT channel type).

comminfo
A communications information object.

listener
A listener.

namelist
A namelist.

process
A process.

queue A queue.

qmgr A queue manager.

service
A service.

topic A topic.

The default value is `all`.

-o *outputType*
Specifies the type of output for the command.

outputType can be one of the following values:

mqsc Multi-line MQSC that can be used as direct input to **runmqsc**

1line MQSC with all attributes on a single line for line diffing

setmqaut
setmqaut statements; valid only when `-x authrec` is specified

grtmqaut
Generates IBM i syntax for granting access to the objects.

The default value is `mqsc`.

-u *userID*
If a userID is specified, a password is requested.

-q Specifies the name of the reply-to queue used when configuration information is retrieved.

-r Specifies the name of the remote queue manager/transmit queue when queued mode is used.

If this parameter is omitted, the configuration for the directly connected queue manager (specified with the **-m** parameter) is dumped.

-m Specifies the name of the queue manager to connect to.

The default value is the default queue manager.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- If any object is not at the default value, the `-a` option must be used if the dumped configuration is used to restore the configuration.
- The **dmpmqcfg** command dumps only subscriptions of type MQSUBTYPE_ADMIN, that is, only subscriptions that are created by using the MQSC command **DEFINE SUB** or its PCF equivalent. The output from **dmpmqcfg** is a runmqsc command to enable the administration subscription to be re-created. Subscriptions that are created by applications by using the MQSUB MQI call of type MQSUBTYPE_API are not part of the queue manager configuration, even if durable, and so are not dumped by **dmpmqcfg**.
- The user must have MQZAO_OUTPUT (+put) authority to access the command input queue (SYSTEM.ADMIN.COMMAND.QUEUE) and MQZAO_DISPLAY (+dsp) authority to access the default model queue (SYSTEM.DEFAULT.MODEL.QUEUE), to be able to create a temporary dynamic queue if the default reply queue is used. The user must also have MQZAO_CONNECT (+connect) and MQZAO_INQUIRE (+inq) authority for the queue manager, and MQZAO_DISPLAY (+dsp) authority for every object that is requested.
- For more information about this command in IBM MQ, see dmpmqcfg in the IBM MQ documentation.

Examples

- The following command dumps the queue manager configuration for a queue manager `QM1`:

```
dmpmqcfg -m QM1
```

dspmq

Display information about queue managers.

Purpose

You can use the **dspmq** command to display the names and details of the queue managers on the IBM MQ Appliance.

Syntax

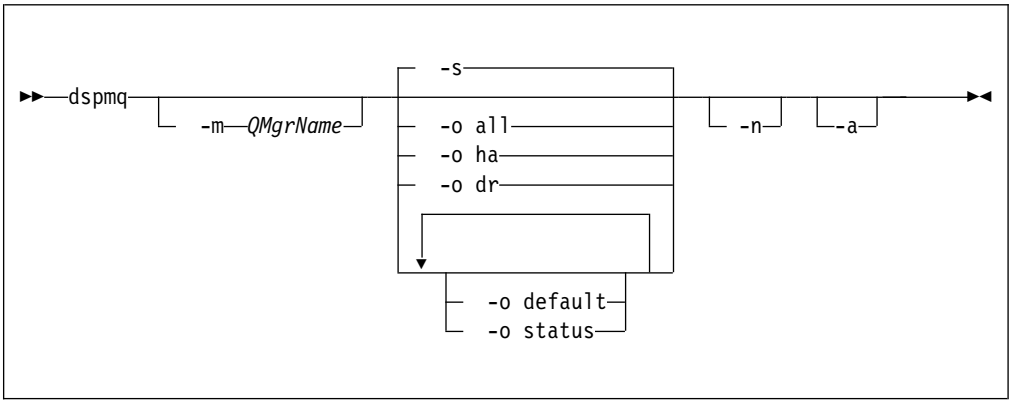

Parameters

-a Specifies that information about only the active queue managers is displayed.

A queue manager is active one or more of the following statements are true:

- The queue manager is running
- A listener for the queue manager is running
- A process is connected to the queue manager

-m *QMgrName*
Specifies which queue manager to display the details for.

If no queue manager name is specified, all queue managers are displayed.

-n Specifies that the translation of output strings is suppressed.

-s
Specifies that the operational status of the queue managers is displayed.

This parameter is the default status setting. It is equivalent to **-o status**.

-o all
Specifies that the operational status of the queue managers is displayed.

-o default
Specifies that the default queue manager status is displayed.

-o ha
Specifies that the HA type is displayed.

-o dr
Specifies that disaster recovery information is displayed. Displays the port that the data replication listener on both appliances uses and the IP address used by the remote appliance.

-o status
Specifies that the operational status of the queue managers is displayed.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- The queue manager can be in any of the following states:

  - Starting
  - Running
  - Running as standby
  - Running elsewhere
  - Quiescing
  - Ending immediately
  - Ending pre-emptively
  - Ended normally
  - Ended immediately
  - Ended unexpectedly
  - Ended pre-emptively
  - Status not available
- For more information about this command in IBM MQ, see dspmq in the IBM MQ documentation.

Examples

- The following command displays queue managers on the appliance:

```
dspmq -o all
```

endmqm

Stop a queue manager.

Purpose

You can use the **endmqm** command to stop a queue manager. This command stops a queue manager in one of three modes:

- Controlled or quiesced shutdown
- Immediate shutdown
- Preemptive shutdown

Syntax

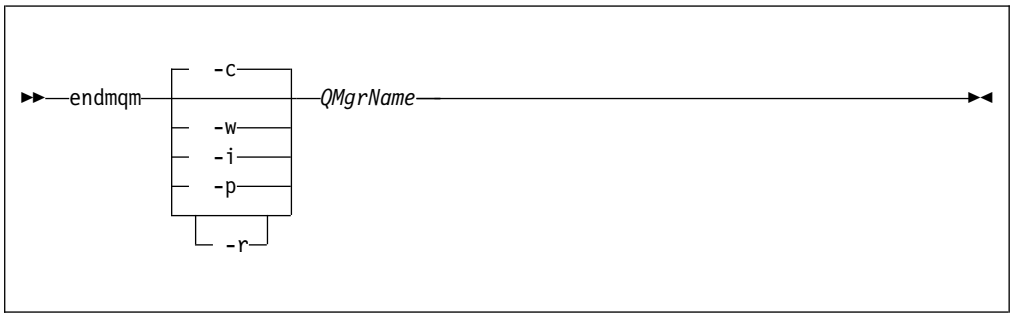

Parameters

QMgrName
Specifies the name of the message queue manager that you want to stop.

This parameter is required.

-c Specifies that the queue manager ends in a controlled (or quiesced) shutdown.

In a controlled shutdown, the queue manager stops after all applications are disconnected. Any MQI calls currently being processed are completed. Control is returned to you immediately and you are not notified of when the queue manager is stopped.

This parameter is the default.

-i Specifies that the queue manager ends in an immediate shutdown.

In an immediate shutdown, the queue manager stops after it all the MQI calls currently being processed are completed. Any MQI requests made after the command starts fail. Any incomplete units of work are rolled back when the queue manager is next started. Control is returned after the queue manager ends.

-p Specifies that the queue manager ends in a preemptive shutdown.

In a preemptive shutdown, the queue manager might stop without waiting for applications to disconnect or for MQI calls to complete. This behavior can give unpredictable results for your applications. Therefore, use this type of shutdown only after other **endmqm** commands fail to stop the queue manager.

-r Specifies that client connectivity can be re-established with other queue managers in their queue manager group.

The client might not reconnect to the same queue manager. Depending on the MQCONNX reconnect option the client uses, and the definition of the queue manager group in the client connection table, the client might reconnect to a different queue manager. You can configure the client to force it to reconnect to the same queue manager.

-w Specifies that the queue manager ends in a wait shutdown.

In a wait shutdown, the queue manager stops after all applications are disconnected. Any MQI calls currently being processed are completed. Control is returned to you after the queue manager stops.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- This command does not affect the attributes of the queue manager.
- The **endmqm** command affects any client application that is connected to the queue manager by a server-connection channel. The effect is equivalent to a STOP CHANNEL command in one of the following modes:
  - If the `-c`, or `-w` parameters are used, the mode is QUIESCE.
  - If the `-i` parameter is used, the mode is FORCE.
  - If the `-p` parameter is used, the mode is TERMINATE.
- If an **dspmq** command is entered in the time between the applications disconnecting and the queue manager stopping, the **dspmq** command might report the status as `Ending immediately`, even if a controlled shutdown was requested.
- For more information about this command in IBM MQ, see endmqm in the IBM MQ documentation.

Examples

- The following command ends the queue manager that is named QM1 in a controlled way:

  ```
  endmqm QM1
  ```

- The following command ends the queue manager that is named QM2 immediately:

  ```
  endmqm -i QM2
  ```

Related commands

- crtmqm (Create queue manager)
- strmqm (Start queue manager)
- dltmqm (Delete queue manager)

runmqsc

Run MQSC commands on a queue manager.

Purpose

You can use the **runmqsc** command to start the runmqsc prompt for a queue manager. From the runmqsc prompt you can directly enter MQSC commands to perform administration tasks. For example, you can define, alter, or delete a local queue.

Syntax

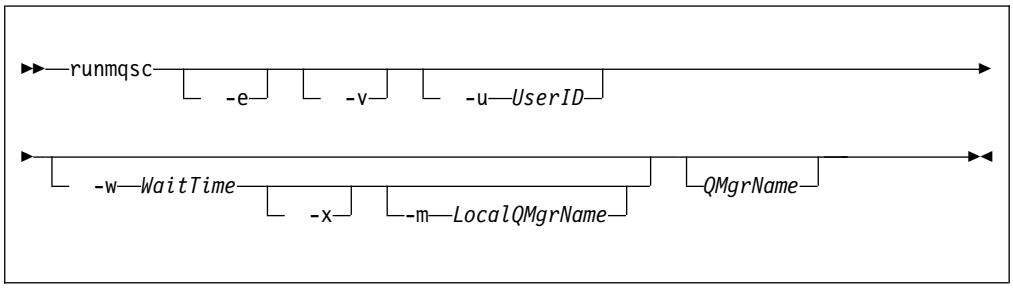

Parameters

-e Specifies that the source text of the MQSC commands is not copied into a report.

This parameter is useful when you enter commands interactively.

-v Specifies that the commands entered are to be verified without performing the action.

You cannot use this parameter with a remote queue manager. That is, the **-w** and **-x** parameters are ignored if specified at the same time as **-v**.

-u *UserID*
Specifies the user ID that the queue manager is accessed with. You are prompted for a matching password.

If you have configured the CONNAUTH AUTHINFO record with CHCKLOCL(REQUIRED) or CHCKLOCL(REQDADM), you must use the **-u** parameter to be able to administer your queue manager with **runmqsc**.

-w *WaitTime*
Specifies that the MQSC commands are to be run on a remote queue manager. The *WaitTime* specifies how many seconds the command waits for replies from the queue manager. Any replies received after this time are discarded, but the MQSC commands still run.

The *WaitTime* must be a value in the range 1 - 999999.

The replies are received on queue SYSTEM.MQSC.REPLY.QUEUE and the outcome is added to the report. This can be defined as either a local queue or a model queue.

You must have the required channel and transmission queues set up for this. See Preparing channels and transmission queues for remote administration in the IBM MQ documentation.
This parameter is ignored if the **-v** parameter is specified.

-x Specifies that the remote queue manager is running under z/OS. The MQSC commands are then written in a form suitable for the z/OS command queue.

This parameter applies only if the **-w** parameter is also specified.

-m *LocalQMgrName*
Specifies the local queue manager that you want to use to submit commands to the remote queue manager.

The default value is the local default queue manager.

This parameter applies only if the **-w** parameter is also specified.

QMgrName
Specifies the name of the target queue manager on which to run the MQSC commands.

The default value is the default queue manager.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- To stop the **runmqsc** command, use the **end** command. You can also use the **exit** or the **quit** command.
- For a full list of MQSC commands and their syntax, see The MQSC commands in the IBM MQ documentation.
- The **runmqsc** command takes its input from `stdin`. You can enter MQSC commands interactively by taking `stdin` from the keyboard. Alternatively, you can enter MQSC commands in a file, and run a sequence of frequently used commands by redirecting the input from the file.
- When the commands are processed, the results and a summary are put into a report that is sent to `stdout`. Therefore, you can redirect the output report to a file.

Examples

- The following command starts the `runmqsc` prompt for the default queue manager:

  ```
  runmqsc
  ```

- The following command starts the `runmqsc` prompt for the queue manager `QM1`:

```
runmqsc QM1
```

From the `runmqsc` prompt you can directly enter MQSC commands.

strmqm

Start a queue manager.

Purpose

You can use the **strmqm** command to start a queue manager.

If the queue manager is part of a high availability configuration, it might start on the other appliance if that is identified as the queue manager's preferred appliance. You can use the status command to check which is the queue manager's preferred appliance. See "status" on page 572.

You can use the **strmqm** command to start a queue manager in the primary role in a disaster recovery configuration. If you use the command to try to start a queue manager in the secondary role, you receive an error message.

Syntax

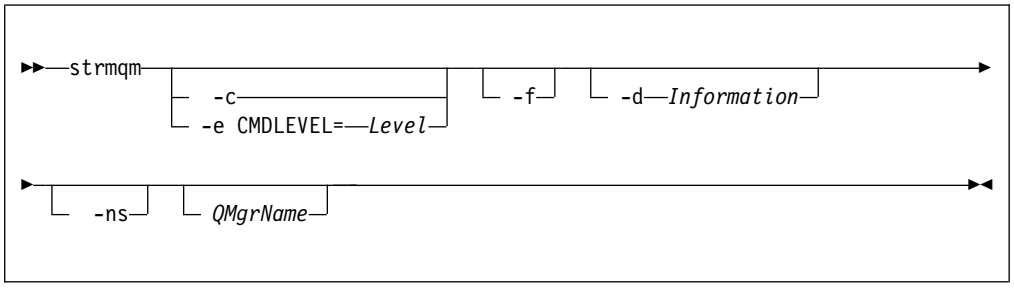

Required parameters

None.

Optional parameters

-c For an HA queue manager, this option only has an effect if it is used with the -ns option.

Specifies that the queue manager default and system objects are to be reset.

Any non-default values for the queue manager default and system objects are replaced with the default values.

The queue manager is stopped after the default and system objects are reset. After you have reset the default and system objects for the queue manager, you must use the **strmqm** command again to start the queue manager.

If you run `mq strmqm -c` on a queue manager that is being used as a IBM MQ Managed File Transfer coordination queue manager, you must rerun the MQSC script that defines the coordination queue manager objects. This script is in a file called *queue_manager_name*`.mqsc`, which is in the IBM MQ Managed File Transfer configuration directory.

-d *Information*
For an HA queue manager, this option only has an effect if it is used with the -ns option.

Specifies whether information messages are displayed.

You can specify one of the following values for *Information*:

all All information messages are displayed.

minimal
The minimal number of information messages are displayed

none No information messages are displayed.

The default value is `all`.

-e CMDLEVEL=*Level*
For an HA queue manager, this option only has an effect if it is used with the -ns option.

Specifies which command level is enabled for the queue manager.

The queue manager is stopped after the command level is set. After you set the command level for the queue manager, you must use the **mq strmqm** command again to start the queue manager.

You must specify a command level that is greater than the current command level of the queue manager and less than or equal to the maximum command level supported by the IBM MQ Appliance.

This flag cannot be specified with `-c`.

-f For an HA queue manager, this option only has an effect if it is used with the -ns option.

Specifies that the queue manager data directory is to be re-created and file permissions are to be reset.

Use this option if you know that a queue manager is not starting because its data directories are missing or corrupted.

If the command is successful, the queue manager starts. If the queue manager fails to start because the configuration information is missing, re-create the configuration information, and restart the queue manager.

You must not use this parameter to re-create the queue manager data directories if you can restore the directories by correcting the configuration. However, you must use this parameter to re-create the queue manager data directory if you are performing media recovery for a queue manager.

-ns
Specifies that the channel initiator, the command server, the listeners, and the services are not started automatically when the queue manager starts. Also specifies that a high availability system does not start. If you start a queue manager that is normally part of an HA configuration, the queue manager starts on the issuing appliance, even if it is not the preferred location. If that appliance subsequently fails, the queue manager will not fail over to the other appliance in the HA group.

QMgrName
Specifies the name of the queue manager to start.

The default value is the default queue manager.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes

mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

- The **strmqm** command is not required to start a newly-created queue manager. An HA queue manager is started automatically on creation.
- The **strmqm** command is required to restart a stopped HA queue manager. In this case, the queue manager is started on the appliance that is the preferred location for the queue manager, regardless of which appliance the command is issued on. If the HA preferred location is not set, the queue manager starts on the same appliance that it stopped on.
- For more information about creating and activating a backup queue manager, see Backing up and restoring WebSphere MQ queue manager data in the IBM MQ documentation.
- For more information about this command in IBM MQ, see strmqm in the IBM MQ documentation.

Examples

- The following command starts the queue manager `QM1`:

```
strmqm QM1
```

Related commands

- crtmqm (Create queue manager)
- endmqm (Start queue manager)
- dltmqm (Delete queue manager)

status

Reports disk usage, CPU usage, and memory usage across the appliance or for a specific queue manager. Also reports additional information for a queue manager running in a high availability configuration, or a disaster recovery configuration.

Purpose

You can use the **status** command to get information about the disk usage, CPU usage, and memory usage for the appliance or for a specific queue manager.

If the queue manager is running in a high availability configuration, the following information is also reported:

- The high availability role of the queue manager (reported as `Primary` or `Secondary`).
- The current high availability status:

  Normal
  The appliances in the disaster recovery configuration are operating normally.

  This appliance in standby mode
  This status means that the appliance has been suspended (by using the **sethagrp -s** command).

  Secondary appliance in standby mode
  This status means that the other appliance in the HA pair has been suspended (by using the **sethagrp -s** command).

Both appliances in standby mode
This status means that both appliances in the HA pair have been suspended (by using the **sethagrp -s** command).

Secondary appliance unavailable
This status means that the connections to the other appliance in the HA pair have been lost.

Partitioned
If the two appliances lose the replication interface, the queue managers are immediately flagged as partitioned. The running queue manager might accumulate out-of-sync data. The other queue manager remains in standby with no out-of-sync data. When the connection is remade, the partitioned state is automatically resolved and replication is resumed.

Synchronization in progress
This status is displayed when the primary queue manager is replicating data to the secondary queue manager.

Inactive
The queue manager is inactive on both appliances in the HA pair.

- The preferred appliance setting for the queue manager, set to `This Appliance` or `Other Appliance`.
- The percentage complete of a synchronization operation. This information is shown only when the status is `Synchronization in progress`.
- The estimated time at which a synchronization will complete. This information is shown only when the status is `Synchronization in progress`.
- The amount of out-of-sync data that exists on this instance of the queue manager. This is the amount of data written to this instance of the queue manager since it entered the partitioned state. This information is shown only when the status is `Partitioned`.

If the queue manager is running in a disaster recovery configuration, the following information is also reported:

- The disaster recovery role of the queue manager (reported as `Primary` or `Secondary`).
- The current disaster recovery status:

Normal
The appliances in the disaster recovery configuration are operating normally.

Synchronization in progress
This status can mean that initial replication is completing, or there has been a failure of the disaster recovery replication network and the queue manager has switched into synchronization mode to catch up as quickly as possible.

Partitioned
Queue manager data on the appliances is out of step, and cannot be automatically resolved. The **makedrprimary** and **makedrsecondary** commands must be used to resolve the situation. When this status is displayed on one of the appliances in a disaster recovery pair, the other appliance might display the **remote appliance unavailable** status, because the connection was lost before it detected the partitioned status.

Remote appliance(s) unavailable
The status means that the connection to the other appliance in the disaster recovery configuration has been lost.

Inactive
The queue manager is in the secondary role on both appliances.

Inconsistent
This status is shown only when the queue manager is in the secondary role and synchronization is in progress. If you use the **makedrprimary** command on a queue manager that is in this state, the queue manager reverts to the snapshot of its data that was taken before it entered the inconsistent state.

Reverting to snapshot
This status is shown when the queue manager is in the secondary role, and the **makedrprimary** command is issued when the queue manager is in the `inconsistent` state. The queue manager is reverted to the current snapshot of its data such that it can run.

Remote appliance(s) not configured
This status is shown when the **crtdrprimary** command has been run, to specify that a queue manager has the primary role, but no **crtdrsecondary** command has been run on the other appliance in the disaster recovery pair.

- The percentage complete of a synchronization operation. This information is shown only when the status is `Synchronization in progress`.
- The estimated time at which a synchronization will complete. This information is shown only when the status is `Synchronization in progress`.
- The amount of out-of-sync data that exists on this instance of the queue manager. This is the amount of data written to this instance of the queue manager since it entered the partitioned state. This information is shown only when the status is `Partitioned`.
- The percentage complete of a reversion to snapshot operation. This information is shown only when the status is `Reverting to snapshot`.

If the queue manager is part of both an HA and a DR configuration, then both HA and DR information is displayed.

Syntax

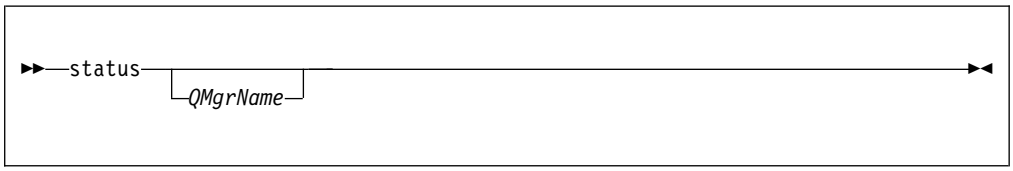

Parameters

QMgrName
Specifies the name of the queue manager for which the status summary is returned.

If this parameter is omitted, a summary of all disk and memory usage on the appliance is returned.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- The information that is returned for the appliance includes the following information:
  - The size and usage of the system memory
  - The CPU usage of the system
  - The size and usage of the internal disk. If the appliance has queue managers that belong to a disaster recovery configuration, the size information includes the space reserved for the snapshot logical volume for each disaster recovery queue manager.
  - The size and usage of the system volume
  - The number of FDCs and the disk space used
  - The disk space used by trace
- The information that is returned for a queue manager includes the following information:
  - The queue manager name
  - The queue manager status
  - The CPU usage of the queue manager
  - The memory usage of the queue manager. If this is a disaster recovery queue manager, this figure does not include the additional memory required for the snapshot image. Note that creating a primary queue manager in a disaster recovery configuration fails if there is insufficient memory bor both the queue manager data, and the snapshot of the queue manager data.
  - The amount of the queue manager file system used by the queue manager
- The information that is returned for a high availability queue manager can also include the following information:
  - The operational state of the HA group
  - The replication status of the queue manager (if synchronization is in progress)
  - The preferred appliance for the queue manager
  - Whether a partitioned situation has been detected, and if it has, the amount of 'out-of-sync' data held
- The information that is returned for a disaster recovery queue manager can also include the following information:
  - The disaster recovery role (primary or secondary)
  - The disaster recovery status
  - The percentage complete if synchronization is in progress
  - The estimated time to completion if synchronization is in progress
  - The amount of out-of-sync data if the disaster recovery system is partitioned
  - The percentage complete if reversion to snapshot is in progress
  - The number of logical writes not yet completed by the primary instance of a queue manager to the secondary instance.
  - The number of logical bytes not yet written by the primary instance of a queue manager to the secondary instance.

Examples

- The following command returns a report for the appliance:
  ```
  status
  ```
- The following command returns a report for a specific queue manager, `QM1`:
  ```
  status QM1
  ```

High availability commands

crthagrp

Creates a high availability (HA) group of appliances.

Purpose

You can use the **crthagrp** command to create an HA group of two appliances. The **prepareha** command must be run on the other appliance before you run **crthagrp**.

You cannot create an HA group if either of the appliances are running queue managers that are configured for disaster recovery.

Syntax

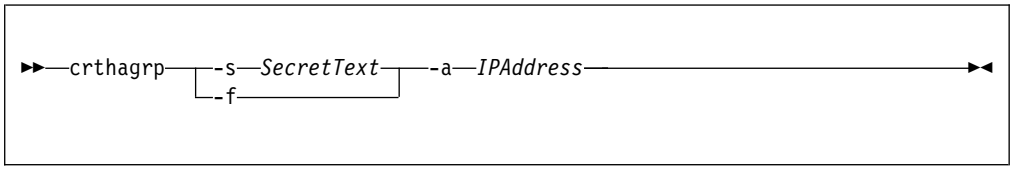

Parameters

-a *IPAddress*
: Specifies the IP address of the HA group primary interface on other appliance in the group. You must specify the IP address using ip v4 dotted decimal notation (for example, "192.0.2.8").

  The IP address specified must be that of the appliance that the command is not run on.

-f
: Forces the creation of an HA group with keys that have already been exchanged between the two appliances. You can only use this command when you have previously created the group using the **-s** parameter to create and exchange the key.

-s *SecretText*
: This argument is used when generating a unique key to be used by the appliances to communicate with one another. Specifies a string that is used to generate a short-lived password. The password is used to set up the unique key for the two appliances. The command **prepareha** must be run first on the other appliance in the HA group, specifying the same `-s` *SecretText* argument.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

- The appliances must be connected to each other with cables inserted in the correct ports. For more information about configuring the appliance hardware for HA, see Configuring the hardware for high availability.

Examples

- The following example shows the creation of an HA group for appliances `appl1` and `appl2` where a new, unique key is generated for communication between the appliances. The HA group primary interface of `appl2` has the IP address `192.0.2.8`, the HA group primary interface of `appl1` has the IP address `192.0.2.7`.

  The following command is run from `appl1`:

  ```
  prepareha -s AGEW1823510HH -a 192.0.2.8
  ```

  The following command is run from `appl2`:

  ```
  crthagrp -s AGEW1823510HH -a 192.0.2.7
  ```

dsphagrp

Displays the status of the appliances in the high availability (HA) group.

Purpose

You can use the **dsphagrp** command to display the status of each appliance in an HA group. The status returned can be `Online`, `Offline`, or `Standby`. The status is `Online` when the appliance is operating normally, `Offline` when some fault has occurred, or `Standby` when the appliance has been suspended by using the **sethagrp -s** command.

Syntax

Parameters

None.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command displays information about the high availability group:

  ```
  dsphagrp
  ```

dlthagrp

Deletes a high availability (HA) group.

Purpose

You can use the **dlthagrp** command to delete an existing HA group.

Syntax

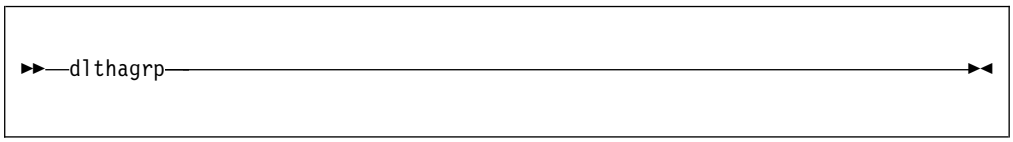

Parameters

None.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- The appliance that you run the command on must have previously been put into an HA group with another appliance. For more information about configuring the appliance hardware for HA, see "Configuring the hardware for high availability" on page 128.
- Before using this command, all HA queue managers must be deleted.
- The HA group is deleted on both appliances in the group. If the other appliance is not available at the time of the delete, the command must be entered on the other appliance to delete the group on that appliance.

Examples

- The following command deletes the HA group that the appliance belongs to:
  ```
  dlthagrp
  ```

makehaprimary

Specifies that an appliance is the 'winner' when resolving a partitioned situation in the high availability group.

Purpose

You use the **makehaprimary** command to specify which appliance in an HA group is considered to have the most up-to-date view of a queue manager after a partitioned situation has occurred.

Syntax

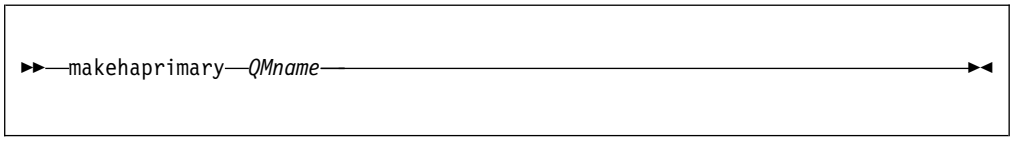

Parameters

QMname
Specifies the name of the queue manager for which the partitioned situation occurred.

The command is run only on the appliance to be identified as the winner.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following example shows the command being run for queue manager QM1.

  The following command is run from the appliance considered to be the winner:

  ```
  makehaprimary QM1
  ```

preparehа

Prepares an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.

Purpose

You use the **prepareha** command to prepare an appliance to be part of an HA group that uses a new, generated key for communication. You run it on the appliance that you do not run **crthagrp** on.

Syntax

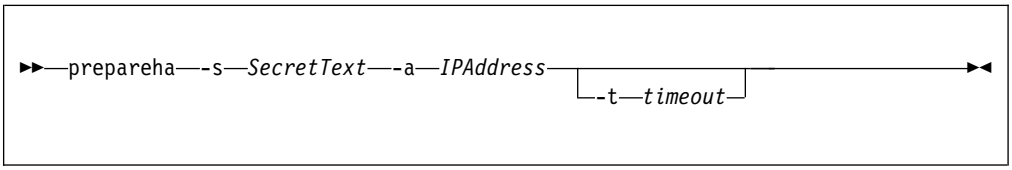

Parameters

-a *IPAddress*
Specifies the IP address of the HA group primary interface on other appliance in the group. You must specify the IP address using ip v4 dotted decimal notation (for example, “192.0.2.8”).

The command is run on only one appliance. The IP address specified must be that of the appliance that the command is not run on.

-s *SecretText*
Specifies a string that is used to generate a short-lived password. The password is used to set up the unique key for the two appliances. After **prepareha** is run, **crthagrp** must be run first on the other appliance in the HA group, specifying the same `-s` *SecretText* argument.

-t *timeout*

Specifies the time period in seconds that you have to run the **crthagrp** command on the other appliance in the group. It defaults to 600 (that is, ten minutes).

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- The appliances must be connected to each other with cables inserted in the correct ports. For more information about configuring the appliance hardware for HA, see "Configuring the hardware for high availability" on page 128.

Examples

- The following example shows the creation of an HA group for appliances `appl1` and `appl2` where a new, unique key is generated for communication between the appliances. The HA group primary interface of `appl2` has the IP address `192.0.2.8`, the HA group primary interface of `appl1` has the IP address `192.0.2.7`.

  The following command is run from `appl1`:

  ```
  prepareha -s AGEW1823510HH -a 192.0.2.8
  ```

  The following command is run from `appl2`:

  ```
  crthagrp -s AGEW1823510HH -a 192.0.2.7
  ```

sethagrp

Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group.

Purpose

You use the **sethagrp** command to pause the high availability (HA) group on an appliance. Any queue managers running on that appliance fail over to the other appliance in the group. You can then use **sethagrp** to resume a previously paused HA group on the appliance.

You can also use the **sethagrp** command to add a standalone queue manager to an HA group, or to remove a queue manager from an HA group and run it as a standalone queue manager.

You cannot remove a queue manager from an HA group if it is also part of a disaster recovery (DR) configuration. In that case, you must remove the DR configuration from the queue manager before you run the **sethagrp** command, see dltdrprimary.

Syntax

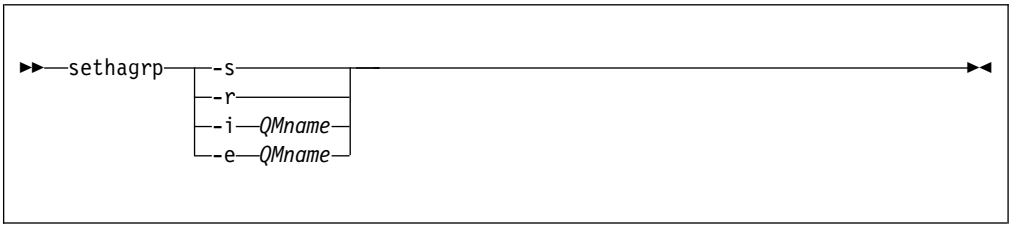

Parameters

-s Suspend the HA group on the appliance into standby mode.

-r Resume the HA group on the appliance from standby mode.

-i *QMname*
Add an existing queue manager to the HA group. The queue manager must not already be under HA control and must be currently stopped. The queue manager is started automatically after the command is completed.

You cannot use this command on a queue manager that is already part of a DR configuration.

-e *QMname*
Remove a queue manager from the HA group. The queue manager must be under HA control and be currently stopped. You must run the command on the appliance that the queue manager was running on when it was stopped. You can discover where the queue manager is running before you stop it by using the **dspmq** command or the **status** *qmanager* command. Either command will report the status as `Running` for the current appliance, or `Running elsewhere` for the other appliance in the HA group.

You cannot use this option if the queue manager is also part of a DR configuration.

Use the **strmqm** command to restart the queue manager after the command is completed.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

sethapreferred

Sets a preferred appliance in the high availability (HA) group for a queue manager to run on.

Purpose

You use the **sethapreferred** command to specify which appliance in an HA pair a queue manager should run on, provided that the appliance is available.

By default, the preferred appliance for a queue manager is the appliance that the queue manager was created on. You can use the **sethapreferred** command in circumstances such as replacing a failed node, or specifying the favored appliance

when an existing queue manager is added to an HA group. The **sethapreferred** is used in conjunction with the **clearhapreferred** command.

You run the command on the appliance that you want to be the preferred appliance, specifying the queue manager name. If the queue manager is currently running on the other appliance, it will fail over to this appliance, provided that is possible (for example, data replication between the two appliances must be up to date).

Syntax

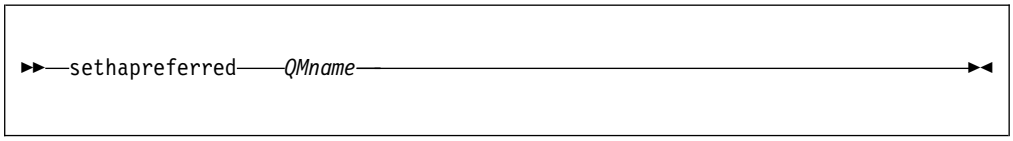

Parameters

QMname
Specify the queue manager that you are setting the preferred appliance for.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

clearhapreferred

Clears the preferred appliance for a queue manager in a high availability (HA) group.

Purpose

You use the **clearhapreferred** command to clear the preferred appliance setting for a queue manager.

By default, the preferred appliance for a queue manager is the appliance that the queue manager was created on. You can use the **clearhapreferred** command to specify that the queue manager has no preferred appliance. You can also use the command when replacing a failed node. The **clearhapreferred** is used in conjunction with the **sethapreferred** command.

Syntax

Parameters

QMname
Specify the queue manager that you are clearing the preferred appliance for. The queue manager must be part of an HA group.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Troubleshooting commands

dltmqras

Deletes the IBM MQ error log files.

Purpose

You can use the **dltmqras** command to periodically purge IBM MQ log files.

For each file deleted, a message in the form `File deleted:` *filename* is written to `MQSystem.log`. You can view `MQSystem.log` by using the **dspmqerr** command without parameters.

Syntax

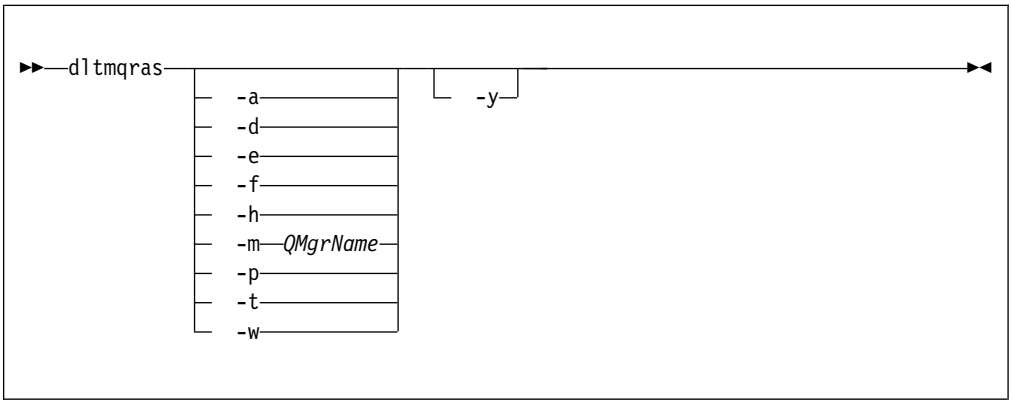

Parameters

-a Specifies that all files of all types apart from queue manager logs are deleted.

-d Specifies that general diagnostics files are deleted.

-e
Specifies that older error logs are deleted. The current error log (`MQSystem.log`) is not deleted.

-f Specifies that FDC files are deleted.

-h Specifies that HA files are deleted.

-m ***QMgrName***
Specifies that service tool output for the specified queue manager are deleted.

-p Specifies that the files in the `mqtemporary:` location are deleted.

-t Specifies that trace files are deleted.

-w
Specifies that console log files are deleted.

-y Specifies that the specified files are deleted without asking for confirmation.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following command purges all types of log files without asking you for confirmation:

  ```
  dltmqras -a -y
  ```

- The following command purges FDC files:

  ```
  dltmqras -f
  ```

- The following command purges all the service tool output for queue manager qm1:

  ```
  dltmqras -m qm1
  ```

dspmqerr

Displays the IBM MQ error log files.

Purpose

You can use the **dspmqerr** command to view or list the IBM MQ error log files.

You can view a list of the files available, then repeat the command specifying a file name to view a specific file. If you specify neither the list argument nor a file name, you can view the default file of the specified type.

The command is based on the UNIX `less` command. The `less` command provides controls for navigating the contents of a file, and you can use these controls when viewing system error logs.

Syntax

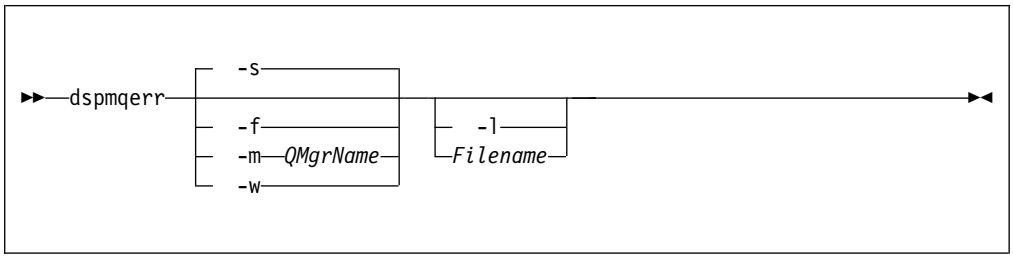

Parameters

-f Specifies that the file type to return is FDC.

-s Specifies that the file type to return is system-wide error log.

-m *QMgrName*
Specifies that the file type to return is the log or logs for the specified queue manager

-w Specifies that the file type to return is an IBM MQ Console log.

-l Lists the files available.

filename
Specifies the particular file to view.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Examples

- The following commands all display the system error log:

```
dspmqerr
dspmqerr -s
dspmqerr MQSystem.log
```

- The following command lists all the error logs on the appliance (but not FDC files):

```
dspmqerr -l
```

- The following command lists all the FDC files:

```
dspmqerr -f -l
```

- The following command lists all the IBM MQ Console files:

```
dspmqerr -w -l
```

- The following command lists all the log files for the queue manager QM1:

```
dspmqerr -m QM1 -l
```

- The following command lists the first log file in the log file directory for the queue manager QM1:

```
dspmqerr -m QM1
```

- The following command lists the log file for the queue manager QM1 named `AMQERR02.LOG`:

```
dspmqerr -m QM1 AMQERR02.LOG
```

- The following command displays the FDC file named `AMQ12345.FDC`:

```
dspmqerr -f AMQ12345.FDC
```

- The following command displays the IBM MQ Console file named `messages.log`:

```
dspmqerr -w messages.log
```

dspmqrte

Determine the route that a message has taken through a queue manager network.

Purpose

You can use the **dspmqrte** to generate a trace-route message and put it into a queue manager network. As the trace-route message travels through the queue manager network, activity information is recorded. When the trace-route message reaches its target queue, the activity information is collected and displayed.

Syntax

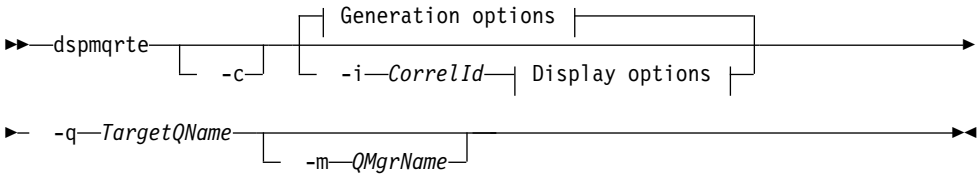

Generation options:

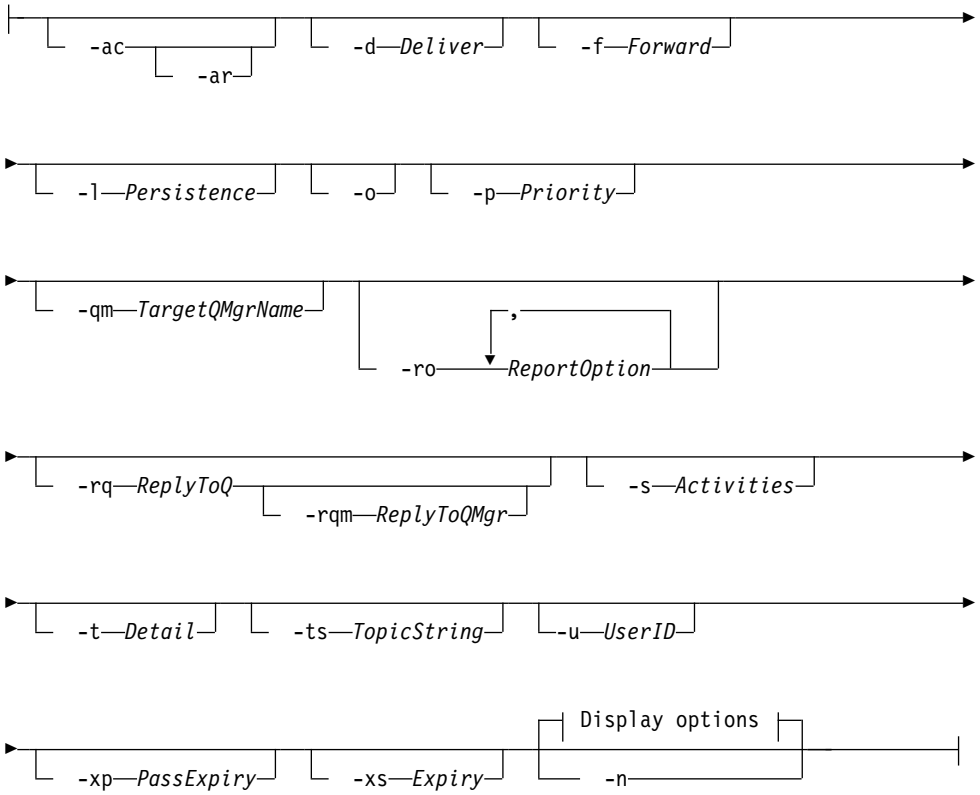

Display options:

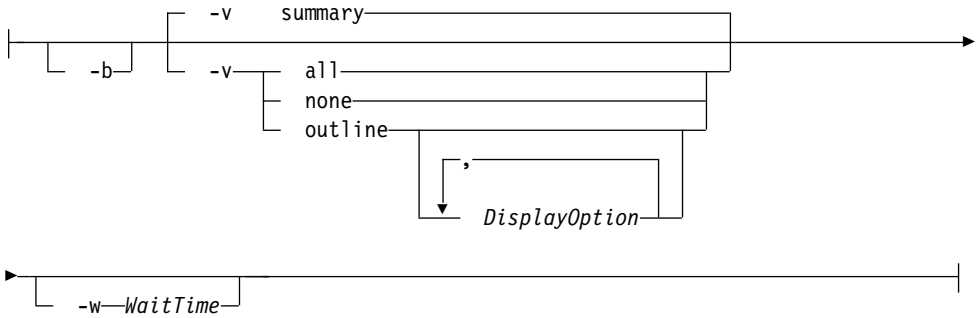

Parameters

-q *TargetQName*
Specifies the name of the target queue to send the trace-route message to.

If the command is being used to view previously gathered activity information, *TargetQName* specifies the name of the queue where the activity information is stored.

This parameter is required.

-c
Specifies that the command connects as a client application.

-i *CorrelId*
Specifies the message identifier of the original trace-route message when displaying previously accumulated activity information.

There can be many activity reports and trace-route reply messages on the queue specified by **-q TargetQName**. Therefore, **-i** can be used to identify the activity reports, or a trace-route reply message, related to a specific trace-route message.

Specify *CorrelId* as a 48 character hexadecimal string.

-m *QMgrName*
Specifies the name of the queue manager to which the command connects.

The name can contain up to 48 characters.

The default value is the default queue manager.

The following parameters are used when the command is used to put a trace-route message into a queue manager network. That is, the parameters are the generation options:

-ac
Specifies that activity information is to be accumulated within the trace-route message.

If you do not specify this parameter, activity information is not accumulated within the trace-route message.

-ar
Specifies that a trace-route reply message containing all accumulated activity information is generated in the following circumstances:

- The trace-route message is discarded by a queue manager.
- The trace-route message is put to a local queue (target queue or dead-letter queue) by a queue manager.
- The number of activities performed on the trace-route message exceeds the value of specified in *-s Activities*.

If you do not specify this parameter, a trace-route reply message is not requested.

-d *Deliver*
Specifies whether the trace-route message is to be delivered to the target queue on arrival.

Deliver can be one of the following values:

yes On arrival, the trace-route message is put to the target queue, even if the queue manager does not support trace-route messaging

no On arrival, the trace-route message is not put to the target queue.

The default value is **no**.

-f *Forward*
Specifies the type of queue manager that the trace-route message can be forwarded to.

Forward can be one of the following values:

all The trace-route message is forwarded to any queue manager.

If forwarded to a queue manager before Version 6.0, the trace-route message is not recognized and can be delivered to a local queue despite the value of the `-d` parameter.

supported
The trace-route message is only forwarded to a queue manager that honors the value of the `-d` parameter.

The default value is `supported`.

-l *Persistence*
Specifies the persistence of the generated trace-route message.

Persistence can be one of the following values:

yes The generated trace-route message is persistent. (MQPER_PERSISTENT)

If you use this value, you must specify the parameter **-rq ReplyToQ**. The reply-to queue must not resolve to a temporary dynamic queue.

no The generated trace-route message is not persistent. (MQPER_NOT_PERSISTENT).

q The generated trace-route message inherits its persistence value from the queue specified by **-q TargetQName**. (MQPER_PERSISTENCE_AS_Q_DEF).

A trace-route reply message, or any report messages, returned shares the same persistence value as the original trace-route message.
The default value is `no`.

-o Specifies that the target queue is not bound to a specific destination.

Typically this parameter is used when the trace-route message is to be put across a cluster. The target queue is opened with option MQOO_BIND_NOT_FIXED.

If you do not specify this parameter, the target queue is bound to a specific destination.

-p *Priority*
Specifies the priority of the trace-route message.

The value of *Priority* is either greater than or equal to 0, or MQPRI_PRIORITY_AS_Q_DEF. MQPRI_PRIORITY_AS_Q_DEF specifies that the priority value is taken from the queue specified by **-q TargetQName**.
The default is that the priority value is taken from the queue specified by **-q TargetQName**.

-qm *TargetQMgrName*
Specifies the target queue manager for the target queue.

The target queue is specified with **-q TargetQName**.

The default is that the queue manager to which the command is connected is used as the reply-to queue manager.

-ro *ReportOption*
ReportOption can be one or more of the following values specified in a comma-separated list:

none Specifies that no report options are set.

activity
The report option MQRO_ACTIVITY is set.

coa The report option MQRO_COA_WITH_FULL_DATA is set.

cod The report option MQRO_COD_WITH_FULL_DATA is set.

exception
The report option MQRO_EXCEPTION_WITH_FULL_DATA is set.

expiration
The report option MQRO_EXPIRATION_WITH_FULL_DATA is set.

discard
The report option MQRO_DISCARD_MSG is set.

The default value is `activity, discard`.

-rq *ReplyToQ*
Specifies the name of the reply-to queue that all responses to the trace-route message are sent to.

If the trace-route message is persistent, or if the `-n` parameter is specified, a reply-to queue must be specified that is not a temporary dynamic queue.

If you do not specify this parameter, the system default model queue, SYSTEM.DEFAULT.MODEL.QUEUE is used as the reply-to queue. Using this model queue causes a temporary dynamic queue to be created.

-rqm *ReplyToQMgr*
Specifies the name of the queue manager where the reply-to queue is located.

The name can contain up to 48 characters.

If you do not specify this parameter, the queue manager to which the command is connected is used as the reply-to queue manager.

-s *Activities*
Specifies the maximum number of recorded activities that can be performed on behalf of the trace-route message before it is discarded.

This parameter prevents the trace-route message from being forwarded indefinitely if caught in an infinite loop.

The value of *Activities* is either greater than or equal to 1, or MQROUTE_UNLIMITED_ACTIVITIES. MQROUTE_UNLIMITED_ACTIVITIES specifies that an unlimited number of activities can be performed on behalf of the trace-route message.

If you do not specify this parameter, an unlimited number of activities can be performed on behalf of the trace-route message.

-t *Detail*
Specifies the activities that are recorded.

Detail can be one of the following values:

low Activities performed by user-defined application are recorded only.

medium
Activities specified in **low** are recorded. Additionally, activities performed by MCAs are recorded.

high Activities specified in **low**, and **medium** are recorded. MCAs do not expose any further activity information at this level of detail. This option is available to user-defined applications that are to expose further activity information only. For example, if a user-defined application determines the route a message takes by considering certain message characteristics, the routing logic can be included with this level of detail.

The default value is `medium`.

-ts *TopicString*
Specifies a topic string to which the command is to publish a trace-route message, and puts the command into topic mode.

In this mode, the command traces all of the messages that result from the publish request.

-u *userID*
User ID to use for connecting to a queue manager.

-xp *PassExpiry*
Specifies whether the report option MQRO_DISCARD_MSG and the remaining expiry time from the trace-route message is passed on to the trace-route reply message.

PassExpiry can be one the following values:

yes The report option MQRO_PASS_DISCARD_AND_EXPIRY is specified in the message descriptor of the trace-route message.

If a trace-route reply message, or activity reports, are generated for the trace-route message, the MQRO_DISCARD_MSG report option (if specified), and the remaining expiry time are passed on.

no The report option MQRO_PASS_DISCARD_AND_EXPIRY is not specified.

If a trace-route reply message is generated for the trace-route message, the discard option and remaining expiry time from the trace-route message are not passed on.

The default value is `yes`.

-xs *Expiry*
Specifies the expiry time for the trace-route message, in seconds.

The default value is 60.

-n Specifies that activity information returned for the trace-route message is not to be displayed.

If this parameter is accompanied by a request for a trace-route reply message (**-ar**), or any of the report generating options (**-ro ReportOption**), then a specific (non-model) reply-to queue must be specified using **-rq ReplyToQ**.

After the trace-route message is put to the specified target queue, a 48 character hexadecimal string is returned containing the message identifier of the trace-route message. The message identifier can be used by the command to display the activity information for the trace-route message at a later time. This can be done using the **-i CorrelId** parameter.

By default, activity report messages are requested.

The following parameters are used when the command is used to display collected activity information. That is, the parameters are the display options:

-b Specifies that the command only browses activity reports or a trace-route reply message related to a message.

This parameter allows activity information to be displayed again at a later time.

If you do not specify this parameter, the command gets activity reports or a trace-route reply message related to a message, and deletes them.

-v summary | all | none | outline *DisplayOption*
The values can be the following values:

summary
The queues that the trace-route message was routed through are displayed.

all All available information is displayed.

none No information is displayed.

outline ***DisplayOption***
Specifies display options for the trace-route message.

DisplayOption can be one or more of the following values, using a comma as a separator:

activity
All non-PCF group parameters in Activity PCF groups are displayed.

identifiers
Values with parameter identifiers MQBACF_MSG_ID or MQBACF_CORREL_ID are displayed.

This value overrides `msgdelta`.

message
All non-PCF group parameters in Message PCF groups are displayed.

When this value is specified, you cannot specify `msgdelta`.

msgdelta
All non-PCF group parameters in Message PCF groups, that have changed since the last operation, are displayed.

When this value is specified, you cannot specify `message`.

operation
All non-PCF group parameters in Operation PCF groups are displayed.

traceroute
All non-PCF group parameters in TraceRoute PCF groups are displayed.

If no values are supplied for *DisplayOption* the application name, the type of each operation, and any operation specific parameters are displayed.

The default value is `summary`.

-w *WaitTime*
Specifies the time, in seconds, that the command waits for activity reports, or a trace-route reply message, to return to the specified reply-to queue.

The default value is the expiry time of the trace-route message, plus 60 seconds.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- For more information about this command in IBM MQ, see dspmqrte in the IBM MQ documentation.

Examples

- The following command puts a trace-route message into a queue manager network with the target queue specified as `TARGET.Q`. Providing queue managers on route are enabled for activity recording, activity reports are generated. Depending on the queue manager attribute, ACTIVREC, activity reports are either delivered to the reply-to queue `ACT.REPORT.REPLY.Q`, or are delivered to a system queue. The trace-route message is discarded on arrival at the target queue.

  ```
  dspmqrte -q TARGET.Q -rq ACT.REPORT.REPLY.Q
  ```

  Providing one or more activity reports are delivered to the reply-to queue, `ACT.REPORT.REPLY.Q`, the command orders and displays the activity information.
- The following command puts a trace-route message into a queue manager network with the target queue specified as `TARGET.Q`. Activity information is accumulated within the trace-route message, but activity reports are not generated. On arrival at the target queue, the trace-route message is discarded. Depending on the value of the target queue manager attribute, ROUTEREC, a trace-route reply message can be generated and delivered to either the reply-to queue, `TRR.REPLY.TO.Q`, or to a system queue.

  ```
  dspmqrte -ac -ar -ro discard -rq TRR.REPLY.TO.Q -q TARGET.Q
  ```

  Providing a trace-route reply message is generated, and delivered to the reply-to queue `TRR.REPLY.TO.Q`, the command orders and displays the activity information that was accumulated in the trace-route message.

endmqtrc

End trace for some or all of the entities that are being traced.

Purpose

You can use the **endmqtrc** command to end tracing for a specified entity, or for all entities.

Syntax

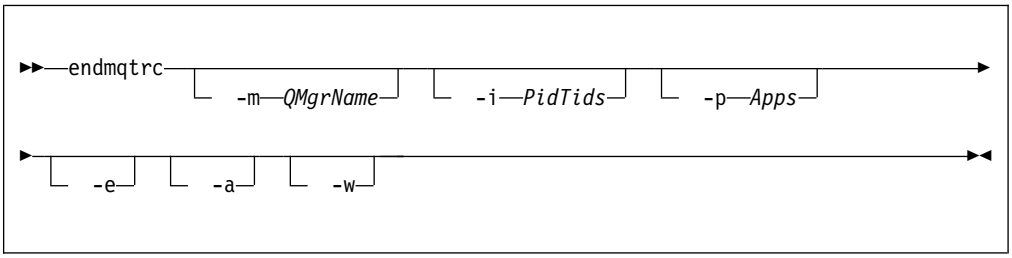

Parameters

-m *QMgrName*
Specifies the name of the queue manager for which to end tracing.

The *QMgrName* supplied must match exactly the *QMgrName* supplied on the **strmqtrc** command. If the **strmqtrc** command used wildcard characters, the **endmqtrc** command must use the same wildcard characters.

A maximum of one `-m` flag can be supplied on the command.

-i *PidTids*
Specifies the process identifier (PID) and thread identifier (TID) for which to end tracing.

You cannot use the `-i` flag with the `-e` flag.

This parameter must be used only under the guidance of IBM Service personnel.

-p *Apps*
Specifies the processes for which to end tracing.

Specify *Apps* as a comma-separated list, with each name in the list specified exactly as the program name would be displayed in the "Program Name" FDC header. You can use an asterisk (*) as a wildcard to match zero or more characters. You can use a question mark (?) to match a single character.

You cannot use the `-p` flag with the `-e` flag.

-e Specifies that early tracing of all processes ends.

You cannot use the `-e` flag with the `-m` flag, the `-i` flag, or the `-p` flag.

-a Ends all tracing.

This flag must be specified alone.

-w Restrict triggering of trace to applications run by an IBM MQ Appliance administrator.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes

mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

- Specifying **endmqtrc** with no parameters has the same effect as specifying **endmqtrc -e**.
- For more information about this command in IBM MQ, see endmqtrc in the IBM MQ documentation.

Examples

- The following command ends tracing of data for a queue manager called `QM1`:

  ```
  endmqtrc -m QM1
  ```

- The following command ends tracing for queue manager `QM2` only. Any other traces that are running are not affected:

  ```
  endmqtrc -m QM2
  ```

Related commands

- "strmqtrc" on page 359

mqrc

Display information about return codes.

Purpose

You can use the **mqrc** command to display information about symbols, return codes, and AMQ messages. You can specify a range of return codes or AMQ messages, or you can specify specific return codes or AMQ messages.

Syntax

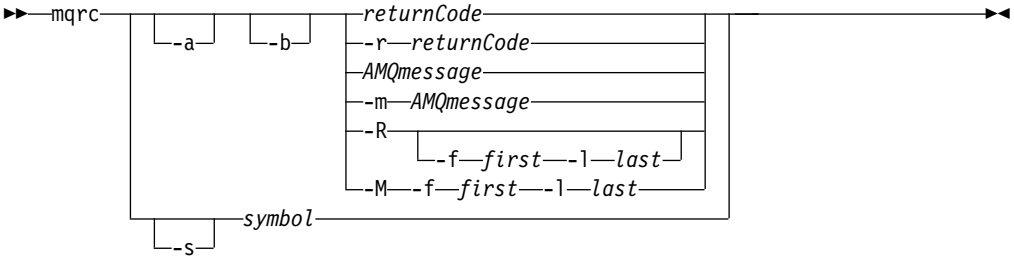

Parameters

returnCode
Specifies the return code to display.

AMQmessage
Specifies the AMQ message to display.

symbol
Specifies the symbol to display.

-a Specifies that all severities are tried to find message text.

-b Specifies that messages are displayed without extended information.

-m *AMQmessage*
Specifies the AMQ message to display.

-M -f *first* -l *last*
Specifies that AMQ messages in a range are displayed from the *first* value to the *last* value.

-r *returnCode*
Specifies the return code to display

-R Specifies that all return codes are displayed.

-R -f *first* -l *last*
Specifies that return codes in a range are displayed from the *first* value to the *last* value.

-s *symbol*
Specifies the symbol to display

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Numeric arguments are interpreted as decimal if they start with a digit 1 - 9, or hex if prefixed with 0x.
- If there is a problem with a message within a range, an indication is displayed before the message text. `?` is displayed if there are no matching return codes for the message. `!` is displayed if the message severity is not the same as the return code severity.
- For more information about this command in IBM MQ, see mqrc in the IBM MQ documentation.

Examples

- This command displays AMQ message 5005:
  ```
  mqrc AMQ5005
  ```
- This command displays return codes in the range 2505 - 2530:
  ```
  mqrc -R -f 2505 -l 2530
  ```

runmqras

Gather diagnostic information together into a single archive to submit to IBM Support.

Purpose

You can use the **runmqras** command to gather diagnostic information from the appliance into a single archive. You can use this command to gather information about an application or appliance failure, possibly for submission into IBM when you report a problem.

By default, the command gathers information such as the FDC files, error logs, product version, and status information. The command does not gather user information that is contained in messages on queues when you use the default setting. However, if you request sections other than default, the data collected might contain user information.

The diagnostic information is written to a zip file named `runmqras_`*`timestamp`*`.zip`, where *timestamp* has the format *yymmdd_HHMMSS*.

The zip file written to the appliance URI mqdiag://, You can retrieve it by using the copy command (see "**copy**" on page 505), or by using the IBM MQ Appliance web UI (see "Managing files by using the IBM MQ Appliance web UI" on page 233). You can also use the **ftp custom** option of the **runmqras** command to copy the trace directly to an FTP server.

Syntax

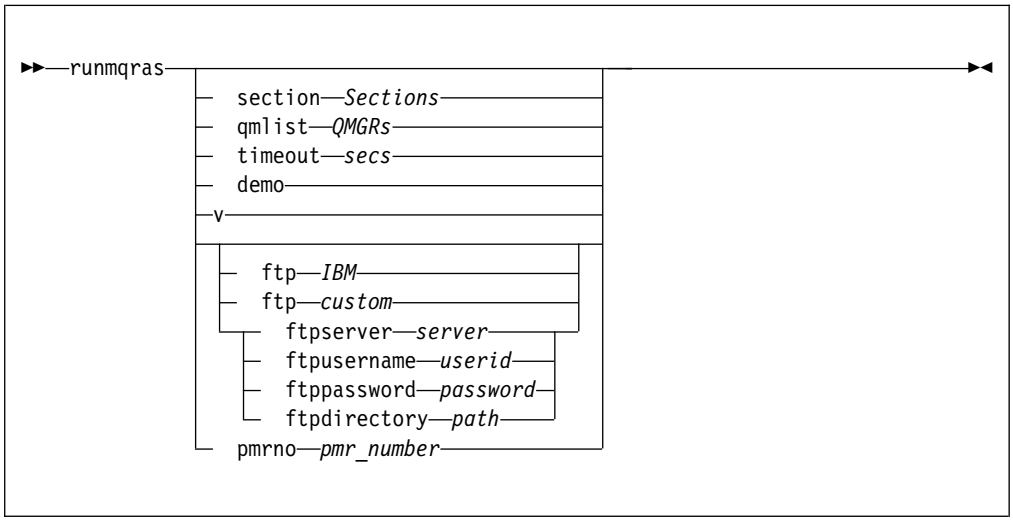

Parameters

section *Sections*
Specifies the optional sections about which to gather more specific information.

By default, a generic section of documentation is collected. Running without requesting more sections is intended as a starting point for general problem diagnosis, but more specific information can be gathered for a specific problem type. You can specify multiple values for *Sections* in a comma-separated list.

IBM support generally supplies you with the sections to use. Example values for *Sections* are the following values:

all Gathers all possible information, including all trace files, and diagnostics for many different types of problems.

This option results in the generation of a very large file, so you must check that the `mqdiag://` directory does not currently contain trace information. If `mqdiag://` does already contain information, you should copy the files off of the appliance, or send them to IBM support, before running **runmqras** with the **all** section.

cluster
Gather information specific for clustering

defs Gather the queue manager definitions and status

nodefault
Prevents the default collections from occurring, but other explicitly requested sections are still collected.

trace Gather all the trace file information plus the default information.

This option results in the generation of a very large file, so you must check that the `mqdiag://` directory does not currently contain trace

information. If `mqdiag://` does already contain information, you should copy the files off of the appliance, or send them to IBM support, before running **runmqras** with the **trace** section.

webui A diagnostics test is run on the IBM MQ Console, and the results written to the archive.

qmlist *QMRs*
Specifies one or more queue manager names on which the command is to be run.

You can specify multiple queue managers in a comma-separated list.

By default, the command is run on all queue managers.

timeout *secs*
Specifies the default timeout to give an individual command before the command stops waiting for completion.

A value of zero means that the command waits indefinitely.

The default value is `10`.

demo
Specifies that the command is run in demonstration mode.

In demonstration mode, no commands are processed, and no files gathered. However, you can see which commands would be processed, and which files would be gathered in the `console.log` file that is generated as part of the output.

-v Specifies verbose output.

ftp IBM pmrno *number*
Specifies that the collected archive is sent through basic FTP to IBM.

Anonymous FTP is used to deliver the archive into the IBM ECuRep server. This process is identical to submitting the file manually by using FTP.

number must specify a valid IBM PMR (problem record number) against which to associate the archive.

ftp custom
Specifies that the collected archive is sent through basic FTP to a site of your choosing.

When you use this parameter, you must specify the following *ftp parameters*:

ftpserver *server*
Specifies an FTP server name to connect to.

ftpusername *userid*
Specifies the user ID to log in to the FTP server with.

ftppassword *password*
Specifies the password to log in to the FTP server with

ftpdirectory *path*
Specifies the directory on the FTP server to place the resulting file into.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes

mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

- For more information about this command in IBM MQ, see runmqras in the IBM MQ documentation.

Examples

- The following command gathers the default documentation from the installation, and all queue managers on the system:

  ```
  runmqras
  ```

- The following command gathers the default documentation from the installation, and sends it directly into IBM to be associated with PMR number 11111,222,333 using the basic FTP capability:

  ```
  runmqras -ftp ibm -pmrno 11111,222,333
  ```

- The following command gathers the default documentation from a machine, plus all trace files, the queue manager definitions, and status for all queue managers on the system:

  ```
  runmqras -section trace,defs
  ```

- The following command copies the information gathered by **runmqras** from the `mqdiag://` directory on the appliance to another location on a system with the IP address 10.10.1.159:

  ```
  (config)# copy mqdiag://runmqras_160818_221406.zip scp://jrb@10.10.1.159//home/user
  ```

strmqtrc

Start trace at a specified level of detail, or report the level of tracing in effect.

Purpose

You can use the **strmqtrc** command to enable tracing. You can specify the tracing that you want:

- You can trace one or more queue managers.
- You can trace one or more processes. The processes can be either part of the product or customer applications that use the IBM MQ API.
- You can trace specific threads within customer applications, either by thread number or by operating system thread number.
- You can trace events. These events can be either the entry or exit from internal functions or the occurrence of a first failure data capture (FDC).
- You can choose from different levels of trace detail.

Trace files are written to the appliance URI mqtrace://, You can retrieve them by using the copy command (see "**copy**" on page 505), or by using the IBM MQ Appliance web UI (see "Managing files by using the IBM MQ Appliance web UI" on page 233)

Syntax

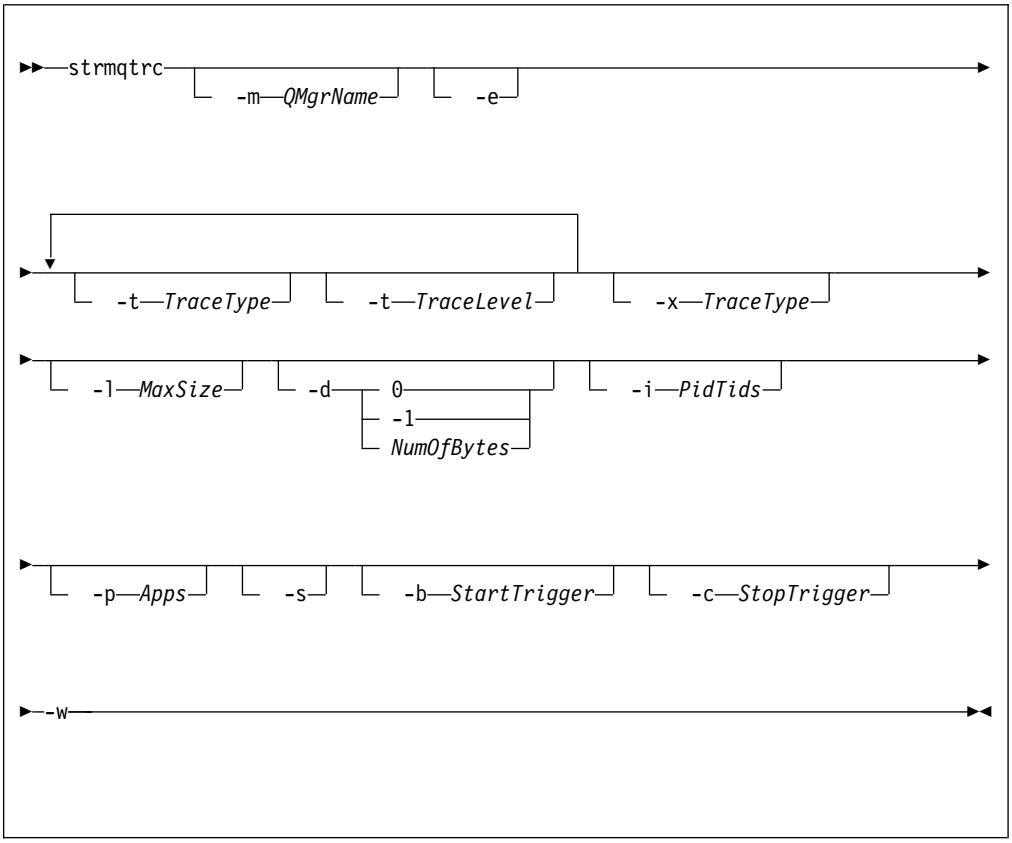

Parameters

-m *QMgrName*
Specifies the name of the queue manager to trace.

You can use an asterisk (*) as a wildcard to replace zero or more characters. You can use a question mark (?) as a wildcard to replace any single character.

-e Specifies that any process that belongs to any component of any queue manager traces its early processing.

You can use this flag to trace the creation or startup of a queue manager.

You cannot use the `-e` flag with the `-m` flag, `-i` flag, the `-p` flag, the `-c` flag, or the `-b` flag.

The default is not to perform early tracing.

-t *TraceType* -t *TraceLevel*
Specifies the points to trace and the amount of trace detail to record.

To specify multiple points to trace, specify multiple `-t` *TraceType* `-t` *TraceLevel* parameters in sequence.

Each *TraceType* can be one of the following values for the points to trace:

all Output data for every trace point in the system. This parameter activates tracing at default detail level.

api Output data for trace points that are associated with the MQI and major queue manager components.

commentary
Output data for trace points that are associated with comments in the components.

comms
Output data for trace points that are associated with data flowing over communications networks.

csdata Output data for trace points that are associated with internal data buffers in common services.

csflows
Output data for trace points that are associated with processing flow in common services.

Explorer
Output data for trace points associated with the IBM MQ Explorer.

Java Output data for trace points associated with applications using the IBM MQ classes for Java™ API.

lqmdata
Output data for trace points that are associated with internal data buffers in the local queue manager.

lqmflows
Output data for trace points that are associated with processing flow in the local queue manager.

otherdata
Output data for trace points that are associated with internal data buffers in other components.

otherflows
Output data for trace points that are associated with processing flow in other components.

parms Activate tracing at default-detail level for flow processing trace points.

remotedata
Output data for trace points that are associated with internal data buffers in the communications component

remoteflows
Output data for trace points that are associated with processing flow in the communications component.

servicedata
Output data for trace points that are associated with internal data buffers in the service component.

serviceflows
Output data for trace points that are associated with processing flow in the service component.

soap Output data for trace points associated with IBM MQ Transport for SOAP.

spldata
Output data for trace points that are associated with buffers and control blocks that use a security policy (AMS) operation.

splflows
Output data for trace points that are associated with entry and exit data for functions that use a security policy (AMS) operation.

ssl Output data that is associated with using GSKit to enable Secure Sockets Layer (SSL) channel security.

versiondata
Output data for trace points that are associated with the version that is running.

The default value is `all`.

Each *TraceLevel* can be one of the following values:

detail Activate tracing at high-detail level for flow processing trace points.

parms Activate tracing at default-detail level for flow processing trace points.

The default value is `parms`.

-x *TraceType*
Specifies the points to exclude from trace.

You can specify the same values for *TraceType* as listed for the `-t` parameter. The default value is `all`.

To specify multiple points to exclude from trace, specify multiple `-x` *TraceType* parameters in sequence.

-l *MaxSize*
Specifies the maximum size of a trace file in megabytes (MB).

The maximum value for *MaxSize* is 2048.

-d 0
Specifies that no user data is traced.

-d -1
Specifies that all user data is traced.

-d *NumOfBytes*
Specifies the number of bytes of data to trace.

For a communication trace, trace the specified number of bytes of data, including the transmission segment header (TSH).

For an MQPUT or MQGET call, trace the specified number of bytes of message data that is held in the message buffer.

Values in the range 1 - 15 are not allowed.

-i *PidTids*
Specifies the process identifier (PID) and thread identifier (TID) to which the trace generation is restricted.

You cannot use the `-i` flag with the `-e` flag.

This parameter must be only used under the guidance of IBM Service personnel.

-p *Apps*
Specifies the named processes to which the trace generation is restricted.

Specify *Apps* as a comma-separated list, with each name in the list specified exactly as the program name would be displayed in the "Program Name" FDC

header. You can use an asterisk (*) as a wildcard to match zero or more characters. You can use a question mark (?) to match a single character.

You cannot use the `-p` flag with the `-e` flag.

-s Specifies that the tracing options that are currently in effect are reported.

You must use this parameter on its own with no other parameters.

-b *Start_Trigger*
Specifies the FDC probe IDs for which tracing must be turned on.

Start_Trigger takes one of the following values:

FDC=*comma-separated list of FDC probe IDs*
Turns tracing on when any FDCs with the specified FDC probe IDs are generated.

You can use an asterisk (*) as a wildcard to match zero or more characters. You can use a question mark (?) to match a single character.

You cannot use the `-b` flag with the `-e` flag.

This parameter must be used only under the guidance of IBM Service personnel.

-c *Stop_Trigger*
Specifies the FDC probe IDs for which tracing must be turned off, or interval in seconds after which tracing must be turned off.

Stop_Trigger takes one of the following values:

FDC=*comma-separated list of FDC probe IDs*
Turns tracing off when any FDCs with the specified FDC probe IDs are generated.

You can use an asterisk (*) as a wildcard to match zero or more characters. You can use a question mark (?) to match a single character.

interval=*n*
Where n is an unsigned integer in the range 1 - 32,000,000.

Turns tracing off n seconds after it starts or, if it tracing is already enabled, turns tracing off n seconds after this instance of the command is entered.

This parameter must be used only under the guidance of IBM Service personnel.

-w Allow any application to trigger trace.

Usage notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Each combination of parameters on an individual invocation of the command are interpreted as having a logical AND between them. You can start the command multiple times, regardless of whether tracing is already enabled. If tracing is already enabled, the trace options that are in effect are modified to those options specified on the most recent invocation of the command.

- Multiple invocations of the command, without an intervening **enqmqtrc** command, are interpreted as having a logical OR between them. The maximum number of concurrent **strmqtrc** commands that can be in effect at one time is 16.
- When a trace file reaches the specified maximum, it is renamed to AMQ*ppppp*.*qq*.TRS and a new AMQ*ppppp*.*qq*.TRC file is started. If a previous copy of an AMQ*ppppp*.*qq*.TRS file exists, it is deleted.
- For more information about this command in IBM MQ, see strmqtrc in the IBM MQ documentation.

Examples

- The following command enables tracing of processing flow from common services and the local queue manager for a queue manager called exampleQM. Trace data is generated at the default level of detail.

  ```
  strmqtrc -m exampleQM -t csflows -t lqmflows -t parms
  ```

- The following command disables tracing of SSL activity on a queue manager called exampleQM. Other trace data is generated at the parms level of detail.

  ```
  strmqtrc -m exampleQM -x ssl -t parms
  ```

- The following command enables high-detail tracing of the processing flow for all components:

  ```
  strmqtrc -t all -t detail
  ```

- The following command enables tracing when FDC KN34650 occurs, and stops tracing when FDC KN346080 occurs. In both cases the FDC must occur on a process that is using queue manager exampleQM:

  ```
  strmqtrc -m exampleQM -b FDC=KN346050 -c FDC=KN346080
  ```

The next examples use the -p and -m flags to show how a combination of parameters on an individual invocation of the command are interpreted as having a logical AND between them. The examples also show how multiple invocations of the command, without an intervening **mq enqmqtrc** command, are interpreted as having a logical OR between them:

1. The following command enables tracing for all threads that result from any executing process that is called amqxxx.exe:

   ```
   strmqtrc -p amqxxx.exe
   ```

2.
   - If you start the following command after the command in step 1, without an intervening **endmqtrc** command, then tracing is limited to all threads that result from any running process that is called amqxxx.exe *and* that are using queue manager exampleQM2:

     ```
     strmqtrc -p amqxxx.exe -m exampleQM2
     ```

   - If you start the following command after the command in step 1, without an intervening **endmqtrc** command, then tracing is limited to all processes and threads that result from running amqxxx.exe *or* that are using queue manager exampleQM2:

     ```
     strmqtrc -m exampleQM2
     ```

Related commands

- "endmqtrc" on page 346

Appliance commands

Use the appliance commands to work with features of the appliance.

The following special characters are used in syntax statements.

[ ] Indicates optional parameters. Parameters that not enclosed in square brackets are required.

| Indicates mutually exclusive parameters. You can use the parameter to the left or the right of the separator. You cannot use all options.

{ } Indicates a set of mutually exclusive parameters when a parameter is required.

Appliance commands

Some of the appliance commands are available at initial log in, before you enter a configuration mode. The other appliance commands are available in most configuration modes.

The following commands are available at initial log in:

Command	Description
"**alias**" on page 437	Creates a command macro.
"**clear intrusion-detected**" on page 439	Clears the detected intrusion when the appliance is in Fail-Safe mode.
"**clock**" on page 440	Sets the date or time.
"**configure terminal**" on page 440	Enters configuration mode.
"**diagnostics**" on page 440	Enters Diagnostics mode. Use this command only at the explicit direction of IBM Support.
"**disable**" on page 441	In privileged mode, this command enters user mode.
"**disconnect**" on page 441	Closes the session for an active user.
"**echo**" on page 442	Echoes text to the console.
"**enable**" on page 442	In user mode, this command enters privileged mode.
"**exec**" on page 442	Calls and runs a configuration script.
"**exit**" on page 443	Applies configuration changes to the running configuration and returns to the parent mode.
"**help**" on page 444	Displays online help.
"**login**" on page 444	Logs a specific account on to the appliance.
"**ntp**" on page 445	Identifies an NTP server.
"**ping**" on page 445	Determines whether the network can reach a remote target.
"**show**" on page 446	Displays configuration or status information
"**shutdown**" on page 447	Restarts or shuts down the appliance.
"**test tcp-connection**" on page 447	Tests the TCP connection to a remote target.

The following commands are available in most configuration modes:

Command	Description
"**admin-state**"	Sets the administrative state for the configuration.
"**cancel**" on page 439	Exits this configuration mode without saving changes to the running configuration.
"**disconnect**" on page 441	Closes the session for an active user.
"**echo**" on page 442	Echoes text to the console
"**exit**" on page 443	Applies configuration changes to the running configuration and returns to the parent mode.
"**help**" on page 444	Displays online help
"**ping**" on page 445	Determines whether the network can reach a remote target.
"**reset**" on page 446	Restores the default values to the configuration.
"**show**" on page 446	Displays configuration or status information
"**test tcp-connection**" on page 447	Tests the TCP connection to a remote target.
"**traceroute**" on page 448	Traces the network path to a target host.

admin-state

This command sets the administrative state for the configuration.

Syntax

Enable the configuration

admin-state enabled

Disable the configuration

admin-state disabled

Parameters

disabled

Sets the configuration to the inactive state.

enabled

Sets the configuration to the active state.

Guidelines

The **admin-state** command sets the administrative state for the configuration. Administrative states are not equivalent to operational states.

- When the administrative state is **enabled**, the operational state can be `up`, `down`, or `pending`.
- When the administrative state is **disabled**, the operational state is always `down`.

Example

Disable the configuration.

```
# admin-state disabled
```

alias

This command creates a command macro.

Syntax

Create a command macro.
alias *alias command*

Delete a command macro.
no alias *alias*

Parameters

alias Specifies the name of the alias.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

command
Specifies a sequence of commands and arguments.

Guidelines

When you use the **alias** command to create a command, you can use one of the following approaches.

- Surround the string in double quotation marks and separate the commands with a semicolon.

```
alias eth0 "configure terminal; interface eth0"
```

- Separate the commands with an escaped semicolon.

```
alias eth0 configure terminal\;interface eth0
```

Use the **no alias** command to delete a command macro.

When you create an alias, use a name that does not conflict with any existing command. For example, if you create the save alias for the **write memory** command, an error occurs when you attempt to use the **save memory report** command. In this case, when you issue the **save memory report** command, the appliance starts the following actions:

- Expands the alias to the **write memory memory report** command.
- Fails invocation and states that the expanded command is unknown.

Examples

- Creates the eth10 alias. When entered, moves to Ethernet configuration mode (with the configure terminal and interface commands) for Ethernet port 10.

```
# alias eth10 configure terminal\;interface eth10
Alias update successful
#
```

- Creates the back2 alias. When entered, moves back two configuration modes. When entered from Validation Credentials mode, moves to Global mode.

```
# alias back2 "exit; exit"
Alias update successful
#
```

- Creates the `update-cfg` alias. When entered, restarts the appliance with an updated configuration script.

```
# alias update-cfg configure terminal\;flash\;del config:///runningconfig.cfg\;
copy http://10.10.1.1/configs/39.3.cfg config:///runningconfig.cfg\;
boot config runningconfig.cfg\;shutdown reload
Alias update successful
#
```

- Deletes the `eth10` alias.

```
# no alias eth10
Alias 'eth10' deleted
#
```

cancel

This command exits this configuration mode without saving changes to the running configuration.

Syntax

`cancel`

Guidelines

The **`cancel`** command exits this configuration mode without saving changes to the running configuration and returns to its parent mode.

clear intrusion-detected

This command clears the detected intrusion when the appliance is in Fail-Safe mode.

Availability

All users, unless your environment enforces RBM on the command line. When RBM enforcement applies to the command line, as defined by the RBM **`apply-cli`** command, this command is available to only the `admin` account (`dp-admin` account on XI50z).

Syntax

`clear intrusion-detected`

Guidelines

The **`clear intrusion-detected`** command clears the detected intrusion when intrusion detection is enabled in System Settings mode. When the appliance detects the intrusion, the appliance starts in Fail-Safe mode. After you clear intrusion detection, you must use the **`shutdown reboot`** command to restart the appliance

Example

Clear intrusion detection of the appliance.

```
(fail-safe)# clear intrusion-detected
Resetting chassis intrusion flag(s) ...
(fail-safe)# shutdown reboot
```

clock

This command sets the date or time.

Syntax

Set the date.

clock *yyyy-mm-dd*

Set the time.

clock *hh:mm:ss*

Parameters

yyyy-mm-dd

Specifies the date in four-digit year, two-digit month, and two-digit day format. When you set the date, separate each value with a hyphen (-).

hh:mm:ss

Specifies the time in two-digit hour, two-digit minute, and two-digit second format. When you set the time, separate each value with a colon (:).

Guidelines

The **clock** command sets the date or time for the appliance. This command is also available in Global mode.

Examples

- Set the date to 8 August 2007.

```
# clock 2007-08-08
Clock update successful
#
```

- Set the time to 8:31 PM.

```
# clock 20:31:00
Clock update successful
#
```

configure terminal

This command enters Global mode.

Syntax

configure terminal

Guidelines

The **configure terminal** command enters Global mode. In this mode, you can create system-wide resources for various system service, configure global behaviors, and enter specialized configuration modes.

diagnostics

This command enters Diagnostics mode. Use this command only at the explicit direction of IBM Support.

Syntax

diagnostics

Guidelines

The **diagnostics** command enters Diagnostics mode. For details about the available commands, see the online help.

Attention: Use this command only at the explicit direction of IBM Support.

disable

In privileged mode, this command enters user mode.

Syntax

disable

Availability

Users who are currently in privileged mode.

Guidelines

The **disable** command exits privileged mode and enters user mode. Privileged and user modes are user-type classes for the appliance.

To exit user mode and enter privileged mode, use the **enable** command.

To close the CLI connection, when not in a global mode, use the **exit** command. When in global mode, the **exit** command returns to initial login mode.

Examples

- Exits privileged mode and enters user mode.

```
# disable
Exiting privileged mode.
>
```

disconnect

This command closes the session for an active user.

Syntax

disconnect *session*

Parameters

session

Specifies the session ID.

Guidelines

The **disconnect** command closes a user session. To list the active user sessions, use the **show users** command.

On XI50z, the dp-admin account cannot disconnect a session that belongs to any other administrative account.

Examples

List the active users and closes the session for the user that is associated with session ID 36.

```
# show users
Session ID  Name ...
...
36          egsmith2 ...
...
# disconnnect 36
Session 36 closed
```

echo

This command echoes text to the console.

Syntax

echo *text*

Parameters

text Specifies the text to echo to the console.

Guidelines

The **echo** command specifies the text to echo to the console.

enable

In user mode, this command enters privileged mode.

Syntax

enable

Availability

Users who are currently in user mode.

Guidelines

After you enter the **enable** command, the CLI prompts for an account name and password. Only authenticated accounts can enter privileged mode. Use privileged mode to provide initial access and to start and to shut down the appliance.

To exit privileged mode and enter user mode, use the **disable** command.

To close the CLI connection, when not in a global mode, use the **exit** command. When in global mode, the **exit** command returns to initial login mode.

Examples

Exit user mode and enter privileged mode.

```
> enable
Username: admin
Password: ********
```

exec

This command calls and runs a configuration script.

Syntax

exec *URL*

Parameters

URL Identifies the location of the configuration file.

- When the file is on the appliance, this parameter takes the *directory*:///*file* format.

  Where:

  directory
  Identifies a local directory. Generally, the directory is `config` or `local`.

  file Identifies the file in the directory.
- When the file is remote and the transport protocol is HTTP or HTTPS, this parameter takes one of the following formats.
  - `http://`*user*`:`*password*`@`*host*`/`*file*
  - `https://`*user*`:`*password*`@`*host*`/`*file*

  The host name can be an IP address or, when DNS is enabled, a qualified host name.

Guidelines

The **exec** command enables modularity of configuration scripts. For example, you can include all service configuration commands in a script called `services.cfg` and all certificates configuration commands in the `cert.cfg` script.

A main configuration script can consist entirely of a series of **exec** commands.

Example

Run the specified configuration scripts.

```
# configure terminal
# exec config:///housekeeping.cfg
# exec config:///interfaces.cfg
# exec config:///crypto.cfg
# exec config:///services.cfg
#
```

exit

This command applies configuration changes to the running configuration and returns to the parent mode.

Syntax

exit

Guidelines

The **exit** command exits the current configuration and applies all changes to the running configuration. To save these changes to the startup configuration, use the **write memory** command.

When you enter the **exit** command from user or privileged mode, this command closes the CLI connection. In all other modes, the command returns to its parent

mode. When issued from the top most parent, the command closes the CLI connection.

Example

Apply all changes made to the `valcred-1` validation credentials configuration . Exit Validation Credentials mode, and returns to Crypto mode. Exit Crypto mode, and returns to Global mode. Persist all configuration changes to the startup configuration. Close the CLI connection.

```
(config valcred valcrec-1)# exit
(config-crypto)# exit
(config)# write memory
(config)# exit
# exit
Goodbye
```

help

This command displays online help.

Syntax

help [*command*]

Parameters

command
Specifies the command name.

Guidelines

The **help** command list the available commands or provides information about a specific command.

- Without arguments, list the commands that are available in the current mode.
- With an argument, displays information about the specific command when that command is available in the current mode.

login

This command logs a specific account on to the appliance.

Syntax

login

Guidelines

After you enter the **login** command, the CLI prompts you for an account name and password.

- The `admin` account, the `dp-admin` account on XI50z, privileged accounts, and group-specific accounts log on to Privileged Mode. The prompt for this mode is the hash (#) character.
- User accounts log on to User Mode. The prompt for this mode is the greater than (>) character.

After an initial login, the CLI prompts you to change the password for the account.

ntp

This command identifies an NTP server.

Syntax

ntp *server* [*interval*]

no ntp

Parameters

server Specifies the IP address or host name.

interval
Specifies the number of seconds between synchronizations with the NTP server. The default value is 900.

Guidelines

The **ntp** command identifies the NTP (Network Time Protocol) server. After you identify an NTP server, the appliance functions as a Simple Network Time Protocol (SNTP) client as described in RFC 2030.

From the CLI, the appliance supports the configuration of only one NTP server. Although the CLI supports only one NTP server, you can use a web management interface to identify multiple NTP servers. When more than one NTP server is identified, the appliance contacts the first NTP server in the list. If this server does not respond, the appliance contacts the next server in the list. If you used a web management interface to identify more than one NTP server, do not use the CLI to modify the NTP service.

Attention: The **ntp** command replaces the entire list with the one identified NTP server.

Use the **no ntp** command to remove the use of NTP servers.

ping

This command determines whether the network can reach a remote target.

Syntax

ping *address*

ping *host* [*IP-version*]

Parameters

address
Specifies the IP address of the target.

host Specifies the host name of the target.

IP-version
Optional: Identifies the IP version to use when resolving an ambiguous host name to an IP address. An ambiguous host name occurs when the DNS publishes an IPv4 address and an IPv6 address. If not specified, resolves to the preferred IP version as defined by the **ip-preference** command in DNS Settings mode.

-4 Identifies the target as an IPv4 host.

-6 Identifies the target as an IPv6 host.

This parameter applies to ambiguous host names only. Although you specify **-6**, the output will show IPv4 output if the DNS published an IPv4 address only. Conversely, if you specify **-4** and the DNS publishes an IPv6 address only, the output will show IPv6 output.

Guidelines

The **ping** command sends 6 Internet Control Message Protocol (ICMP) echo-request messages to the specified host with a one second interval between each message and reports the results.

reset

This command restores the default values to the configuration.

Syntax

reset

Guidelines

The **reset** command restores mode-specific properties to their default values. Properties that lack default values, are unchanged.

This command has no effect on properties that are specific to MQCLI mode.

Default values that are assigned by the **reset** command are not applied until you use the **exit** command to save changes and exit the mode.

Example

Restore default values to the configuration and returns to the parent mode.

```
# reset
# exit
#
```

show

This command displays configuration or status information

Syntax

show [*argument*]

Parameters

argument
Specifies the specific configuration or status provider.

Guidelines

The **show** command displays configuration information or status information that is relevant to the provided argument. Without an argument, the result differs depending on where you entered the command.

- Within the initial login, lists available arguments.

- Within a configuration mode, lists the current configuration.

shutdown

This command restarts or shuts down the appliance.

Syntax

shutdown

shutdown reboot [*seconds*]

shutdown halt [*seconds*]

Parameters

reboot Shuts down and restarts the appliance.

halt Shuts down the appliance without restarting. The power to the appliance remains on.

seconds
Specifies the number of seconds before the appliance starts the shutdown operation. Enter a value in the range 0 - 65535. The default value is 10.

Guidelines

The **shutdown** command shuts down, or shuts down and restarts the appliance. Without parameters, the command restarts the appliance.

The appliance restarts with the startup configuration that is specified by the **boot config** command and the startup firmware image that is specified by the **boot image** command. Without a designated startup configuration or firmware image, the appliance restarts with the configuration and firmware image that were active when you issued the **shutdown** command.

Examples

- Wait 10 seconds to shut down and restart the appliance.

```
# shutdown reboot
Reboot in 10 second(s).
#
```

- Wait 1 minute to shut down the appliance.

```
# shutdown halt 60
Shutdown in 60 second(s).
#
```

test tcp-connection

This command tests the TCP connection to a remote target.

Syntax

Test the connection with an IP address.
test tcp-connection *address port* [*seconds*]

Test the connection with a host name.
test tcp-connection *host port* [*seconds*] [*IP-version*]

Parameters

address
Specifies the IP address of the target.

host Specifies the host name of the target.

IP-version
Optional: Identifies the IP version to use when resolving an ambiguous host name to an IP address. An ambiguous host name occurs when the DNS publishes an IPv4 address and an IPv6 address. If not specified, resolves to the preferred IP version as defined by the **ip-preference** command in DNS Settings mode.

-4 Identifies the target as an IPv4 host.

-6 Identifies the target as an IPv6 host.

This parameter applies to ambiguous host names only. Although you specify `-6`, the output will show IPv4 output if the DNS published an IPv4 address only. Conversely, if you specify `-4` and the DNS publishes an IPv6 address only, the output will show IPv6 output.

port Specifies the target port.

seconds
Specifies the number of seconds to wait for a response from the target. The default value is 10.

Guidelines

The **test tcp-connection** command verifies TCP connectivity from the appliance to a specific target.

Examples

- Test the TCP connection to the specified host on port 80 with the default timeout value of 10 seconds.

```
# test tcp-connection ragnarok 80
TCP connection successful
#
```

- Test the TCP connection to the specified IP address on port 21. The timeout value is 5 seconds.

```
# test tcp-connection 192.168.77.27 21 5
TCP connection successful
#
```

traceroute

This command traces the network path to a target host.

Syntax

Trace the path by IP address.
traceroute *address*

Trace the path by host name.
traceroute *host* [*IP-version*]

Parameters

address
Specifies the IP address of the target.

host Specifies the host name of the target.

IP-version
Optional: Identifies the IP version to use when resolving an ambiguous host name to an IP address. An ambiguous host name occurs when the DNS publishes an IPv4 address and an IPv6 address. If not specified, resolves to the preferred IP version as defined by the **ip-preference** command in DNS Settings mode.

-4 Identifies the target as an IPv4 host.

-6 Identifies the target as an IPv6 host.

This parameter applies to ambiguous host names only. Although you specify `-6`, the output will show IPv4 output if the DNS published an IPv4 address only. Conversely, if you specify `-4` and the DNS publishes an IPv6 address only, the output will show IPv6 output.

Guidelines

The **traceroute** command traces the route that packets actually take to their target host. The output shows the IP address of the hops (for example, gateway or routers) and the round trip time.

The **traceroute** command is intended for use in network testing, measurement, and management. In other words, use this command for manual fault isolation. Because of the load it imposes on the network, do not use this command during typical operations or from automated scripts.

While the **ping** command confirms IP network reachability, you cannot pinpoint and improve some isolated problems. Consider the following situation:

- When there are many hops (for example, gateways or routers) between the appliance and the target host, and there seems to be a problem somewhere along the path. The target host could have a problem, but you need to know where a packet is actually lost.
- The **ping** command hangs up and provides no reason for a lost packet.

The **traceroute** command can inform you where the packet is located and why the route is lost. If your packets must pass through routers and links, which belong to and are managed by other organizations or companies, it is difficult to check related routers. The **traceroute** command provides a supplemental role to the **ping** command.

Example

Confirm an available TCP connection to `loki`.

```
# traceroute loki
```

Appliance user commands

You can use the appliance user commands to configure appliance users on the IBM MQ Appliance.

The appliance user commands can be run from the command line interface in user configuration mode. To enter user configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:

```
config
```

2. From global configuration mode, enter user configuration mode:

   ```
   user name
   ```

   Where *name* identifies the user that you want to configure. If you are creating a new user, *name* can contain up to 128 characters. The following characters are valid:

   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name comprising a single period, or including two periods together, is not permitted)

3. Type `exit` to save your changes and leave user configuration mode, then type `exit` again to leave global configuration mode.

access-level

This command assigns an account type of an account.

Syntax

access-level { **group-defined** | **privileged** | **user** }

Parameters

group-defined
Specifies that the user is granted access privileges as defined by a specific User Group. The user must later be assigned to this group with the `group` command.

privileged
Assigns executive access to the account. A privileged account has virtually the same access levels as the `admin` account. It differs only in that a privileged account cannot delete the `admin` account.

user Deprecated - The **user** keyword is deprecated. Instead, assign the user account to a group-defined account type and define access restrictions through the group.

Guidelines

By default, newly created accounts are assigned the `user` access level.

group

This command associates an account with a user group.

Syntax

group *name*

Parameters

name Specifies the name of a user group.

Guidelines

The **group** command associates an account with a user group.

password

This command manages the password for an account.

Syntax

password *password*

Parameters

password

Specifies the password for the account. A password can contain only printable characters and must be 5 - 20 characters in length.

Guidelines

The **password** command assigns a password to a new account, or changes the password of an existing account.

You must assign when you create an account.

Appliance user group commands

You can use the appliance user group commands to configure appliance user groups on the IBM MQ Appliance.

The appliance user group commands can be run from the command line interface in user group configuration mode. To enter user group configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:
   ```
   config
   ```
2. From global configuration mode, enter user configuration mode:
   ```
   usergroup name
   ```

   Where *name* identifies the user group that you want to configure. If you are creating a new user group, *name* can contain up to 128 characters. The following characters are valid:
   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name comprising a single period, or including two periods together, is not permitted)
3. Type `exit` to save your changes and leave user group configuration mode, then type `exit` again to leave global configuration mode.

access-policy

This command assigns an access policy.

Syntax

access-policy "*statement*"

Parameters

statement

Specifies the policy statement to add. A policy statement takes the following form:

address/*domain*/*resource*?[Name=*name*]&Access=*permission* [&*field*=*value*]

address An IP address or host alias for a local interface (Ethernet or VLAN) on the appliance. The special value * matches all appliance addresses.

domain The name of an application domain. This policy applies to only resources in the identified domain.

- The special value * matches all domains.
- A PCRE can match select domains.

resource

The resource type to which this policy applies. The special value * matches all resource types.

Name=*name*

Optional: Identifies by name an instance of the specified resource type. You can use a PCRE; for example, `foo.*` to specify all resources that start with `foo`.

Access=*permission*

The permission string assigns permissions. The string is cumulative and connected by plus (+) signs. For example, the string `a+d+x+r+w` represents add, delete, execute, read, and write permissions.

field=*value*

Optional: The field token must be one of the additional fields that can be added to the string. The corresponding value can be a PCRE.

Guidelines

The **access-policy** command assigns one or more access policy statements to the user group. If there are more than one statement, the statements are cumulative. If more than one statement applies to the same resource, the most specific statement applies. For example, given the following two statements any member of this user group can read all objects but has complete access privileges to the web management interface:

```
*/*/*?Access=r
*/*/mgmt/web-mgmt?Access=r+w+a+d+x
```

It is not possible to remove a specific access policy from the CLI. If you run the **no access-policy** command, all access policies are removed. To remove a specific access policy from a user group, use the GUI.

Examples

- Add full access privileges to all resources and read only access for GUI login and network interface resources to members of the `appdev` user group.

```
# usergroup appdev
User group configuration mode
# access-policy "*/*/*?Access=r+w+a+d"
# access-policy "*/*/login/web-mgmt?Access=r"
# access-policy "*/*/network/interface?Access=r"
# exit
Usergroup update successful
#
```

add

This command adds a command suite.

Syntax

add *name*

Parameters

name Specifies the name of the command suite.

Guidelines

The **add** command adds a command suite to the user group. To display a list of available command groups, use the **add** command, without arguments.

Editing of a user group with the **add** or **delete** command does not affect the access privileges of existing user accounts that are assigned to the user group. The updated access privileges that are defined by the edited user group are assigned to user accounts that are later added to the user group.

To apply the updated access privileges to existing user accounts, delete the account with the **no user** command and then re-create the account with the **user** command.

Example

Add access to the `configuration`, `urlmap`, `urlrefresh`, `matching`, and `stylesheetpolicy` command suites to members of the `stylesheet` user group.

```
# usergroup stylesheets
User group configuration mode
# add configuration
# add urlmap
# add urlrefresh
# add matching
# add stylesheetpolicy
# exit
Usergroup update successful
#
```

delete

This command removes a command suite.

Syntax

delete *name*

Parameters

name Specifies the name of the command suite.

Guidelines

The **delete** command removes a command suite from this user group. To display a list of available command groups, enter the **delete** command without arguments.

Editing of a user group with the **add** or **delete** command does not affect the access privileges of existing user accounts that are assigned to the user group. The updated access privileges that are defined by the edited user group are assigned to user accounts that are later added to the user group.

To apply the updated access privileges to existing user accounts, delete the account with the **no user** command and then re-create the account with the **user** command.

Example

Remove access to the `configuration` command suite for the `stylesheets` user group.

```
# usergroup stylesheets
User group configuration mode
# delete configuration
# exit
Usergroup update successful
#
```

Crypto commands

You can use the crypto commands to manage certificates on the IBM MQ Appliance.

The crypto commands can be run from the command line interface in crypto configuration mode. To enter crypto configuration mode, type `crypto`.

Note: For certificate commands for queue managers, see "Queue manager certificate management commands" on page 374

authcookie-cache-delete

This command deletes AuthCookie cache entries.

Syntax

authcookie-cache-delete *key*

Parameters

key

Specifies the hashed key for the cache entry.

Guidelines

The **authcookie-cache-delete** command deletes AuthCookie entries from the cache.

To view all entries in the cache, use the **show authcookie-cache-status** command.

cert-monitor

This command enters Certificate Monitor mode.

Syntax

cert-monitor

Guidelines

The **cert-monitor** command enters Certificate Monitor mode. The certificate monitor scans the expiration date of all certificates.

certificate

This command creates an alias for an X.509 certificate.

Syntax

certificate *alias URL* [**password** *password*] [**ignore-expiration**]

certificate *alias URL* [**password-alias** *password-alias*] [**ignore-expiration**]

no certificate *alias*

Parameters

alias Specifies the alias for the certificate.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods. For a z/OS® certificate, see your z/OS Communications Server documentation.

URL Specifies a URL that identifies the file that contains the certificate.

- If stored in the public cryptographic area, takes the pubcert:///*filename* form.
- If stored in the private cryptographic area, takes the *filename* form.
- If retrieved from z/OS, takes the saf-cert://*nssclient*/*filename* form.

password ***password***
Specifies the plaintext password to access the certificate file.

password-alias ***password-alias***
Specifies the alias for the encrypted password to access the certificate file.

ignore-expiration
Specifies an optional keyword to allow the creation of a certificate before its activation date (the NotBefore value in the certificate) or after its expiration date (the NotAfter value in the certificate). Although the certificate is in the up operational state, any configuration that references the certificate uses the internal expiration values.

In other words, the certificate itself is in the `up` operational state, but validation credentials, firewall credentials, or identification credentials that reference the certificate adhere to the internal expiration values.

- If the certificate is for certificate chain validation from validation credentials and the certificate is invalid, validation fails.
- If the certificate is for certificate chain validation from identification credentials, the appliance sends the certificate to the SSL peer for a connection. The peer can reject the certificate as invalid.

Guidelines

The **`certificate`** command creates an alias for an X.509 certificate.

The `password` or `password-alias` keyword is required only when a certificate file is password-protected.

To use the `password-alias` keyword, you must have created an alias. Use the `password-map` command to create the password alias.

Use the **`certificate`** command with the **`key`** and **`idcred`** commands to create identification credentials. Identification credentials consist of a certificate, which contains a public key, and the corresponding private key.

Use the **`certificate`** command with the **`valcred`** command to create validation credentials. Validation credentials can be used, but are not required, during the SSL handshake to authenticate the certificate from the remote SSL peer.

Use the **`no certificate`** command to delete only the alias for the certificate. The file that contains the certificate material remains on the appliance.

Examples

- Create the `bob` alias for the `bob.pem` X.509 certificate. Store the target certificate in the public cryptographic area.

```
# certificate
bob pubcert:bob.pem
Creating certificate 'bob'
```

- Create the `bob` alias for the `bob.pem` certificate. Store the target certificate in the public cryptographic area. Allow the certificate to be accessed with the `pikesville` plaintext password.

```
# certificate bob pubcert:bob.pem
password pikesville
Creating certificate 'bob'
```

- Create the `bob` alias for the `bob.pem` certificate. Store the target certificate in the public cryptographic area. Allow the certificate to be accessed with the `dundaulk` encrypted password alias.

```
# certificate bob pubcert:bob.pem
password-alias dundaulk
Creating certificate 'bob'
```

- Create the `zicsfCert5` alias for the z/OS `ICSFCERT5` certificate. Use the `nssclient` NSS client to connect to z/OS to retrieve the target certificate from z/OS. Store the target certificate in memory.

```
# certificate zicsfCert5 saf-cert://nssclient/ICSFCERT5
Creating certificate 'zicsfCert5'
```

- Delete the `bob` certificate alias.

```
# no certificate bob
Certificate 'bob' deleted
```

convert-certificate

This command converts a certificate alias to a specific output format and writes it to a file.

Syntax

convert-certificate *alias file* [*format*]

Parameters

alias Specifies the name of the certificate alias.

file Specifies the output file name. Use the `temporary:///mycert.pub` format.

format
Specifies the format for the output file. The supported format is `openssh-pubkey`.

Guidelines

The **convert-certificate** command converts a certificate alias to a specific output format and writes it to a file.

The `openssh-pubkey` format can be used in OpenSSH `authorized_keys` files.

convert-key

This command converts a private key alias to a specific output format and writes it to a file.

Syntax

convert-key *alias file* [*format*]

Parameters

alias Specifies the name of the key alias.

file Specifies the output file name. Use the `temporary:///mykey.pub` format.

format
Specifies the format for the output file. The supported format is `openssh-pubkey`.

Guidelines

The **convert-key**command converts a private key alias to a specific output format and writes it to a file. If the output format includes private fields of the key, the file must be in the same directory as the configured file of the private key alias.

The `openssh-pubkey` format can be used in OpenSSH `authorized_keys` files. The format does not contain any private fields. It contains only public fields.

cookie-attribute-policy

This command enters Cookie Attribute Policy mode.

Syntax

cookie-attribute-policy *name*

no cookie-attribute-policy *name*

Parameters

name Specifies the name of the configuration.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

Guidelines

The **cookie-attribute-policy** command enters Cookie Attribute Policy mode. In this mode, you can define the following cookie attributes to include in a cookie:

- Domain
- Path
- Secure
- HttpOnly
- Max-Age
- Expires
- Custom attributes

To delete a cookie attribute policy, use the **no cookie-attribute-policy** command.

crl

This command enters CRL mode to create or modify a CRL update policy.

Syntax

crl *name* { **http** | **ldap** }

no crl *name*

Parameters

name Specifies the name of the CRL update policy.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)

- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

http Indicates an HTTP-enabled CRL update policy.

ldap Indicates an LDAP-enabled CRL update policy.

Guidelines

The **crl** command enters CRL mode to create or modify a CRL (certificate revocation list) update policy.

- Use the **fetch-url** and **refresh** commands to define an HTTP-enabled CRL update policy
- Use the **bind-dn**, **bind-pass**, **read-dn**, **refresh**, and **remote-address** commands to define an LDAP-enabled CRL update policy.

Use the **no crl** command to delete a CRL update policy.

Examples

- Create the HTTP30 HTTP-enabled CRL update policy.

```
# crl HTTP30 http
Entering CRL mode for 'HTTP30'
```

- Create the LDAP1440 LDAP-enabled CRL update policy.

```
# crl LDAP1440 ldap
Entering CRL mode for 'LDAP1440'
```

- Delete the LDAP1440 LDAP-enabled CRL update policy.

```
# no crl LDAP1440
```

crypto-export

This command creates an export package that contains a certificate or key.

Syntax

Exports a certificate

crypto-export cert *name* **output** *file*

Exports a key on an HSM-enabled appliance

crypto-export key *name* **output** *file* **mechanism hsmkwk**

Parameters

key *name*

Identifies the name of the key.

cert *name*

Identifies the name of the certificate.

output *file*

Identifies the name and location to store the export.

mechanism hsmkwk

Indicates that the export package is exported with an HSM-generated key-wrapping key.

Guidelines

The **crypto-export** command creates an export package that contains a certificate or key.

Example

Create the exportBob.xml export package in the temporary: directory. The package contains the bob certificate.

```
# crypto-export cert bob output temporary:///exportBob.xml
```

crypto-hw-disable

This command schedules how many aspects of the cryptographic card to disable for the next firmware reload.

Syntax

crypto-hw-disable { **none** | **all** | **rsa** }

Parameters

none Indicates that no aspects of the cryptographic card are disabled after the next firmware reload.

all Indicates that all aspects of the cryptographic card are disabled after the next firmware reload. The cryptographic card is disabled completely.

rsa Indicates that only the RSA functionality of the cryptographic card is disabled after the next firmware reload. The random number generator of the cryptographic card can still be used.

Guidelines

The **crypto-hw-disable** command schedules how many aspects of the cryptographic card to disable. The change takes effect at the next firmware reload. This command is independent of the appliance-wide cryptographic mode that is controlled with the **crypto-mode-set** command. If you never set the cryptographic hardware disablement with this command, no aspects of the cryptographic hardware are disabled.

To display the status of the cryptographic card disablement, use the **show crypto-hw-disable** command.

Example

Disable the cryptographic card completely during the next firmware reload.

```
# crypto-hw-disable all
```

crypto-import

This command imports an export package that contains a certificate or key.

Syntax

Imports a certificate

crypto-import cert *alias* **input** *file*

Imports a key on an HSM-enabled appliance

crypto-import key *name* **input** *file* [**password-alias** *alias*]

crypto-import key *name* **input** *file* [**password** *password*]

Parameters

key *name*
Identifies the name of the key.

cert *name*
Identifies the name of the certificate.

input *file*
Identifies the name and location of the stored export package.

password *password*
Optional: Specifies the password that was used to encrypt the input file. This parameter is mutually exclusive to the `password-alias` parameter.

password-alias alias
Optional: Specifies the password that was used to encrypt the input file. This parameter is mutually exclusive to the `password` parameter.

Guidelines

The **crypto-import** command is used in two scenarios to import to an HSM.

- Importing an existing private key not from an HSM
- Importing a key that is exported from an HSM with the key-wrapping key

In the former case, the key might be password-protected, so the import requires the password parameter. In the latter case, the incoming key file is wrapped with the key-wrapping key and a password is never used. The key-wrapping key must be on cloned on the appliance with the **hsm-clone-kwk** command.

Examples

- Import the `exportBob.xml` export package from the `temporary` directory. The package contains the `bob` certificate.

  ```
  # crypto-import cert bob input temporary:///exportBob.xml
  ```

- Import the key file not from an HSM, `bobKey.pem`, to the HSM as `bobKey`. The key file is password-protected. The password is provided as an argument. The imported key can later be exported from the HSM with the HSM key-wrapping key method.

  ```
  # crypto-import key bobKey input temporary:///bobKey.pem password pass1word2
  ```

- Take the `exportBobHSMKey.xml` package that was exported from another HSM with the key-wrapping key and import the `bobHSMKey` key to the HSM. No password argument is used to import a key that was exported from an HSM.

  ```
  # crypto-import key bobHSMKey input temporary:///exportBobHSMKey.xml
  ```

crypto-mode-set

This command sets the appliance-wide cryptographic mode for the next firmware reload.

Syntax

crypto-mode-set *mode*

Parameters

mode Indicates which cryptographic mode to enable. The following keywords are available to indicate the modes to enable:

permissive
Runs the firmware in permissive mode.

fips-140-2-l1
Runs the firmware in FIPS 140-2 Level 1 mode.

Guidelines

The **crypto-mode-set** command sets the cryptographic mode for the appliance. The change takes effect at the next firmware reload. The specified mode remains effective until the command is called with a different mode and the firmware is reloaded. If you never set the cryptographic mode with this command, the appliance runs in permissive mode.

- Use FIPS 140-2 Level 1 mode when you must comply with FIPS requirements.
- Use permissive mode to switch back to normal operations.

When you set the mode to FIPS 140-2 Level 1 mode, you must understand the following:

- FIPS 140-2 Level 1 mode removes support in the firmware's main task for MD2, MD4, MD5, RIPEMD160, single DES, RC2, RC4, Blowfish, and CAST because these algorithms are prohibited by the corresponding specification. These algorithms are only available in the firmware's main task in permissive mode.
- FIPS 140-2 Level 1 mode prohibits the use of public keys smaller than 1024 bits.
- FIPS 140-2 Level 1 mode makes the firmware's main task use a pseudorandom number generator that is compliant with NIST SP800-131a and FIPS 140-2.
- Any cryptographic hardware cards that are not validated to FIPS 140-2 will have their RSA functionality disabled when the appliance is in FIPS 140-2 Level 1 mode.

Use the **show crypto-mode** command to display the status of cryptographic modes.

Examples

- Enable FIPS 140-2 Level 1 mode at the next reload of the firmware.

```
# crypto-mode-set fips-140-2-l1
```

- Change the cryptographic mode back to permissive mode at the next reload of the firmware.

```
# crypto-mode-set permissive
```

fwcred

This command enters Firewall Credentials mode.

Syntax

fwcred *name*

no fwcred *name*

Parameters

name Specifies the name of the configuration.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

Guidelines

The **fwcred** command enters Firewall Credentials mode to manage a firewall credentials list. While in Firewall Credentials mode, use the **key**, **sskey**, and **certificate**. The **key** and **sskey** commands add specified keys to the firewall credentials list. The **certificate** command adds certificates to the firewall credentials list.

The assignment of firewall credentials is optional. Firewall credentials provide a way to specify which keys and certificates the various cryptographic extension functions support for firewall activities. Without firewall credentials, all keys and certificates on the appliance are available to support firewall activities.

To delete a firewall credentials configuration, use the **no fwcred** command.

Examples

- Enter Firewall Credentials mode to create the FWCred-1 firewall credentials.

```
# fwcred FWCred-1
Entering Firewall Credentials mode for 'FWCred-1'
```

- Delete the FWCred-1 firewall credentials.

```
# no fwcred FWCred-1
Firewall Credentials 'FWCred-1' deleted
```

idcred

This command creates identification credentials.

Syntax

idcred *name key-alias certificate-alias* [**ca** *certificate-alias-n*]

Parameters

name Specifies the name of the configuration.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

key-alias
Specifies an existing alias for the private key that is referenced by the identification credentials.

certificate-alias
Specifies an existing alias for the certificate that is referenced by the identification credentials.

ca *certificate-alias-n*
Optional: Identifies an intermediate certificate that is required to establish a chain-of-trust starting with the certificate that is referenced by the certificate-alias and a CA trusted by the remote SSL peer. The list can contain up to 10 intermediate certificates.

Guidelines

The **idcred** command creates identification credentials. An SSL proxy profile uses identification credentials to authenticate itself to a remote peer.

The SSL standard requires an SSL server to authenticate itself to a remote SSL client. An SSL proxy profile operating as an SSL server (in either `reverse` or `two-way` mode) must be assigned identification credentials with which to authenticate itself to a remote SSL client.

The SSL standard allows an SSL server to authenticate the remote client peer. An SSL proxy profile operating as an SSL client (in either `forward` or `two-way` mode) can be assigned a set of identification credentials if the remote SSL server requires authentication. While SSL servers typically do not require client identification, certain highly secure websites can impose such a requirement.

Before you create identification credentials, you must use the following procedure.

1. Use the **key** command to create an alias for the private key.
2. Use the **certificate** command to create an alias for the certificate.

The **no idcred** command deletes only the alias for the identification credentials. The aliases that created the identification credentials and the files that contain the actual certificate and private key remain available for use.

Examples

- Create the `bob` identification credentials that consist of the private key that is aliased by `bob` and the X.509 certificate aliased by `bob`.

```
# idcred bob bob bob
Creating identification credentials 'bob'
```

- Create the `bob` identification credentials that consist of the private key that is aliased by `bob` and the X.509 certificates aliased by `bob` and `bob-intermediate`.

```
# idcred bob bob bob ca bob-intermediate
Creating identification credentials 'bob'
```

- Delete the identification credentials alias `bob`.

```
# no idcred bob
Identification Credentials 'bob' deleted
```

key

This command creates an alias for a private key.

Syntax

key *alias URL* [**password** *password*]

key *alias URL* [**password-alias** *password-alias*]

no key *alias*

Parameters

alias Specifies the alias for the private key.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods. For a z/OS key, see your z/OS Communications Server for details on label names.

URL Specifies a URL that identifies the file that contains the private key.

- To store the private key in the private cryptographic area, the URL takes the *filename* form.
- To store the private key in the public cryptographic area, the URL takes the pubcert:///*filename* form.

  Note: Do not store private key files in the pubcert: directory. This directory is intended for the storage of public certificate files.
- To retrieve the private key from z/OS, the URL takes the saf-key://*nssclient*/*filename* form. A saf-key:// must be a SAF key that is not stored in ICSF.
- To access the remote private key on the z/OS system, the URL takes the saf-remote-key://*nssclient*/*filename* form. A saf-remote-key:// must be a SAF key that is stored in ICSF.

password ***password***
Optional: Identifies the plaintext password that is required to access the private key file.

password-alias ***password-alias***
Optional: Identifies the alias for the encrypted password that is required to access the private key file.

Guidelines

The password or password-alias keyword is required only when a key file is password-protected.

To use the password-alias keyword, you must have created an alias. Use the password-map command to create the password alias.

Use the **key** command with the **certificate** and **idcred** commands to create identification credentials that consist of a certificate, which contains a public key and the corresponding private key.

Use the **no key** command to delete only the alias for the private key. The file that contains the key material remains on the appliance.

Examples

- Create the bob alias for the K2.pem private key. The target key is in the private cryptographic storage area.

```
# key bob K2.pem
Creating key 'bob'
```

- Create the bob alias for the K2.der private key. The target key is in the private cryptographic area and is accessed with the annapolis plaintext password.

```
# key bob K2.der password annapolis
Creating key 'bob'
```

- Create the bob alias for the K2.der private key. The target key is in the private cryptographic area and is accessed with the towson encrypted password alias.

```
# key bob K2.der password-alias towson
Creating key 'bob'
```

- Create the zCert_key alias for the z/OS CERT private key. Use the nssclient NSS client to connect to and retrieve the target key. Cache the target key on the appliance.

```
# key zCert_key saf-key://nssclient/CERT
Creating certificate 'zCert_key'
```

- Create the zicsfCert2_key alias for the z/OS ICSFCERT2 private key. Use the nssclient NSS client to connect to and access the ICSFCERT2 private key but does not retrieve or store the z/OS private key on the appliance.

```
# key zicsfCert2_key saf-remote-key://nssclient/ICSFCERT2
Creating certificate 'zicsfCert2_key'
```

- Delete the bob private key alias.

```
# no key bob
Key 'bob' deleted
```

keygen

This command generates a public-private key pair and a CSR (certificate signing request) for a server.

Syntax

Generates a key pair

> **keygen** [{**C** | **countryName**} *iso-code*] [{**L** | **localityName**} *locality*] [{**ST** | **stateOrProvinceName**} *state*] [{**O** | **organizationName**} *org*] [{**OU** | **organizationalUnitName**} *unit-name*] {**CN** | **commonName**} *server-name* **rsa** {**1024** | **2048** | **4096**} [**gen-object**] [**object-name** *name*] [**gen-sscert**] [**days** *number-days*] [**file-name** *name*] [**export-key**] [**export-sscert**] [**password** *plaintext*] [**password-alias** *alias*] [**using-key** *name*]

Parameters

{C | countryName} *ISO-code*

> Optionally specifies the ISO two-character country identifier for the CSR.

{L | localityName} *locality*
Optionally specifies the city or town name for the CSR. Use a text string up to 64 characters in length. If the string contains spaces, enclose in double quotation marks.

{ST | stateOrProvinceName} *state*
Optionally specifies the unabbreviated state or province name for the CSR. Use a text string up to 64 characters in length. If the string contains spaces, enclose in double quotation marks.

{O | organizationName} *organization*
Optionally specifies the organization name for the CSR. Use a text string up to 64 characters in length. If the string contains spaces, enclose in double quotation marks.

{OU | organizationalUnitName} *unit-name*
Optionally specifies the organizational unit name for the CSR. Use a text string up to 64 characters in length. If the string contains spaces, enclose in double quotation marks.

{CN | commonName} *server-name*
Specifies the fully qualified domain name of the server for the CSR. Use a text string up to 64 characters in length.

rsa {1024 | 2048 | 4096}
Indicates the length of the generated RSA key. The default value is 1024. The generation of a 4096-bit key can take up to 30 seconds.

gen-object
Creates a key management object. To create a certificate management object use the `gen-sscert` property.

object-name *name*
Optionally specifies the names for the objects that are created by the `gen-object` property. If not specified, the value for the `commonName` property is used.

gen-sscert
Optionally creates a self-signed certificate in addition to the private key and CSR.

days *number-days*
Optionally specifies the validity period in days for the self-signed certificate. The default value is 365 days.

file-name *name*
Optionally specifies a common prefix for the generated private key, CSR, and self-signed certificate. If not specified, the value for the `object-name` property is used.

export-key
Optionally creates a copy of the private key in the `temporary:` directory in addition to the one in the `cert:` directory.

export-sscert
Optionally creates a copy of the self-signed certificate in the `temporary:` directory in addition to the one in the `cert:` directory.

password *plaintext*
Optionally specifies the password to encrypt the private key when it is saved to a file.

password-alias *alias*
Optionally specifies a password alias in an existing password map file. This alias is used to decrypt the password.

using-key name
Optionally specifies an existing key object to sign the CSR and any self-signed certificate that is generated. The point of this parameter is to reissue a new CSR or self-signed certificate with the existing key material to do the signature.

Guidelines

CA policies can vary with regard to the amount of information that is required in the CSR. Check with the CA before generating the CSR to ensure that you provide sufficient information.

The `password` or `password-alias` keyword is required only when a key file is password protected.

To use the `password-alias` keyword, you must have created an alias. Use the **password-map** command to create the password alias.

Examples

- Generate a private key and CSR for the specified server. Default conditions apply as follows.
  - The private key (1024 bits in length) is saved as `cert:sample-privkey.pem`.
  - The CSR is saved as `temporary:sample.csr`.
  - The private key file is not password protected

```
# keygen C au L "South Melbourne" ST Victoria
 O "DataPower Australia, Ltd." OU "Customer
 Support" CN www.bob.datapower.com.au
```

- Generate a private key and CSR for the specified server with the following options.
  - The private key (2048 bits in length) is saved as `cert:bob-privkey.pem`.
  - The CSR is saved as `temporary:bob.csr`.
  - The private key file is password protected with the plaintext password `didgeridoo`.

```
# keygen C au L "South
 Melbourne" ST Victoria
 O "DataPower Australia, Ltd." OU "Customer
 Support" CN www.bob.datapower.com.au rsa 2048 out bob password
 didgeridoo
```

- Create a new password map and generate a host key to encrypt the plaintext password `didgeridoo`, and associate the alias `WaltzingMatilda` with the encrypted password. Generate a private key and CSR for the specified server with the following options.
  - The private key (2048 bits in length) is saved as `cert:bob-privkey.pem`.
  - The CSR is saved as `temporary:bob.csr`.
  - The private key file is password protected with the encrypted password `didgeridoo`.

```
# password-map
Please enter alias-name and plaintext password pairs
- Leading and trailing white space is removed
- Enter a blank alias name to finish
Alias-name: WaltzingMatilda
```

```
Plaintext password: didgeridoo
Alias-name:
SSL: password-map saved
# keygen C au
 L "South Melbourne" ST Victoria
 O "DataPower Australia, Ltd." OU "Customer Support"
 CN www.bob.datapower.com.au rsa 2048 out bob
 password-alias WaltzingMatilda
```

password-map

This command manages the encrypted passwords to a password aliases in a password map file.

Syntax

Interactively add an entry to the password map file.

password-map

Delete an entry from the password map file.

delete password-map *alias*

Delete the password map file.

no password-map

Parameters

alias The alias is the reference to a password.

Guidelines

The **password-map** command maps the encrypted password to a password alias in a password map file.

The password map and the locally generated key are saved to separate files on the appliance. Plaintext passwords are not saved on the appliance. Password maps are typically used to protect key and certificate files.

- In commands that use plaintext, or unencrypted passwords, the `password` argument is used to open and read the corresponding file.
- In commands that use encrypted passwords, the `password-alias` argument is the search criteria for the password map file to identify its associated encrypted password. Then the encrypted password is decrypted with the locally generated host key to yield the plaintext password. This password is used to open and read the corresponding file.

An attempt to reference an encrypted password that is not found in the password map results in command failure.

The **password-map** command interactively prompts for `alias:password` pairs.

alias Specifies the name of the alias. This name must consist of alphanumeric characters and cannot contain white space. The length is limited to 127 characters.

password
Specifies the plaintext password. This password must consist of alphanumeric characters but can contain white space (spaces or tabs). Leading and trailing white space is ignored. The length is limited to 127 characters.

You must ensure that synchronization is maintained between the startup configuration and the password map file. You must use the **password-map** command to generate and encrypt aliases for certificate or key passwords before the **certificate** or **key** commands can access files that are protected by an encrypted password. An attempt to reference an encrypted password that is not in the password map results in failure.

Deletion of the password map and host key file has no immediate effect on keys and certificates that are in memory. At restart, however, **key** and **certificate** commands that contain references to aliases in the deleted password map fail unless a new password map was created with the same aliases.

Note: The **password-map** command cannot be used in a configuration script. When found, the command is ignored.

Use the **no password-map** command to delete the password map and host key files.

Examples

- Create a password map and generate the host key to encrypt the two plaintext passwords.

```
# password-map
Please enter alias-name and plaintext password pairs
- Enter a blank alias name to finish
Alias-name: towson
Plaintext password: ********
Re-enter plaintext password: ********
Alias-name: dundaulk
Plaintext password: ********
Re-enter plaintext password: ********
Alias-name:
Password-map saved (2 entries)
```

- Confirm the creation of the password map.

```
# show password-map
2 password-map aliases
  towson
  dundaulk
```

- Add another alias-password pair to the password map.

```
# password-map
A password-map already exists, overwrite? Yes/No [y/n]: n
Appending to current password map...
Please enter alias-name and plaintext password pairs
- Leading and trailing white space is removed
Alias-name: columbia
Plaintext password: ********
Re-enter plaintext password: ********
Alias-name:
Password-map saved (3 entries)
```

- Delete the entry associated with the `columbia` alias.

```
# delete password-map columbia
Deleted passwrod-map alias 'columbia'
password-map saved : 2 entry(s)
```

- Delete the password map.

```
# no password-map
Are you sure you want to remove the password-map? Yes/No [y/n]: y
Deleted saved password-map
```

profile (deprecated)

This command creates a cryptographic profile that specifies an SSL service level.

Syntax

profile *name idCred* [**ssl** *name*] [**ciphers** *cipher-string*] [**option-string** *options-mask*] [**clientcalist** {**on** | **off**}]

profile *name* **%none%** [**ssl** *name*] [**ciphers** *cipher-string*] [**option-string** *options-mask*] [**clientcalist** {**on** | **off**}]

no profile *name*

Parameters

name Specifies the name of the configuration.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

idCred Specifies the name of identification credentials that were created with the **idcred** command. The assignment of identification credentials is required when the cryptographic profile is used with an SSL server, which is required by the SSL specification to authenticate itself to clients. The assignment of identification credentials is optional when the cryptographic profile is used with an SSL client.

%none% Indicates that no identification credentials are assigned to the cryptographic profile.

ssl ***name***
Specifies the name of validation credentials that were created with the **valcred** command.

ciphers ***cipher-string***
Specifies a list of symmetric key-encryption algorithms that are supported by this cryptographic profile. Table 33 list the available keywords. The default value is HIGH:MEDIUM:!aNULL:!eNULL:@STRENGTH.

Table 33. Available algorithm keywords for the cipher string

Algorithm keyword	Meaning
DEFAULT	All cipher suites (including weak export ciphers) except the aNULL and eNULL ciphers.
ALL	All cipher suites except the eNULL ciphers, which must be explicitly enabled.
HIGH	All *high* encryption cipher suites, which are currently AES or 3DES cipher suites.
MEDIUM	All *medium* encryption cipher suites, which are currently 128-bit RC2 and 128-bit RC4 cipher suites.
LOW	All *low* encryption cipher suites, which are currently 64-bit RC4 and 56-bit DES non-export cipher suites.

Table 33. Available algorithm keywords for the cipher string (continued)

Algorithm keyword	Meaning
EXP or EXPORT	Export (weak) encryption algorithms that are eligible for export out of the United States, currently 40 and 56-bit algorithms.
EXPORT40	40-bit export (weak) encryption algorithms.
EXPORT56	56-bit export (weak) encryption algorithms.
eNULL or NULL	NULL ciphers offer no encryption at all and are a security risk. These cipher suites are disabled unless explicitly included.
aNULL	The cipher suites that offer no authentication, which is the anonymous DH algorithms. These cipher suites are vulnerable to *man-in-the-middle* attacks. Use is normally discouraged.
kRSA and RSA	Cipher suites that use RSA key exchange.
kEDH	Cipher suites that use ephemeral DH key agreement. These cipher suites are only supported when the cryptographic profile acts as a client.
aRSA	Cipher suites that use RSA authentication where the certificates carry RSA keys.
aDSS and DSS	Cipher suites that use DSS authentication where the certificates carry DSS keys.
TLSv1	TLS version 1.0 cipher suites.
SSLv3	SSL version 3.0 cipher suites.
SSLv2	SSL version 2.0 cipher suites.
DH	Cipher suites that use DH, including anonymous DH.
ADH	Anonymous DH cipher suites.
3DES	Cipher suites that use triple DES.
DES	Cipher suites that use DES instead of triple DES.
RC4	Cipher suites that use RC4.
RC2	Cipher suites that use RC2.
MD5	Cipher suites that use MD5.
SHA1 or SHA	Cipher suites that use SHA-1.
AES	Cipher suites that use AES.

The cipher string consists of one or more cipher keywords that are separated by colons. Commas or spaces are acceptable separators, but colons are the norm.

The cipher string can take different forms.

- A single cipher suite, such as RC4-SHA.
- A list of cipher suites that contains a certain algorithm, or cipher suites of a certain type. For example, SHA1 represents all ciphers suites that use the SHA-1 digest algorithm.
- A combination of single cipher string that uses the + character, which is used as a logical AND operation. For example, SHA1+DES represents all cipher suites that contain the SHA-1 and the DES algorithms.

Optionally, each cipher keyword can be preceded by the following characters:

! Permanently deletes the cipher from the list. Even if you explicitly add the cipher to the list, it can never reappear in the list.

- Deletes the cipher from the list. You can add this cipher again.

+ Moves the cipher to the end of the list. The + character moves existing ciphers. It does not add them.

If none of these characters is present, the string is interpreted as a list of ciphers to be appended to the current list. If the list includes a cipher that is already in the list, that cipher is ignored. That is, existing ciphers are not moved to the end of the list.

Additionally, the cipher string can contain the `@STRENGTH` keyword at any point to sort the cipher list in order of encryption algorithm key length.

option-string *options-mask*
Optional: Enables various SSL options for the cryptographic profile. Use the string that identifies specific supported SSL options.Table 34 lists the available options.

Table 34. SSL options as string

String value	Description
`OpenSSL-default`	The default value.
`Disable-SSLv2`[1]	Disallows the use of SSL version 2.
`Disable-SSLv3`	Disallows the use of SSL version 3.
`Disable-TLSv1`	Disallows the use of TLS version 1.0.
`Enable-Legacy-Renegotiation`	Allows SSL and TLS renegotiation, which is vulnerable to a man-in-the-middle (MITM) attack that is documented in CVE-2009-3555.
`Enable-Compression`	Allows compression.
`Disable-TLSv1d1`	Disallows the use of TLS version 1.1.
`Disable-TLSv1d2`	Disallows the use of TLS version 1.2.
[1] SSL protocol version 2 is deprecated.	

When you use the string value, use a + character to join values. For example, to disallow both SSL version 3 and TLS version 1.0, enter `Disable-SSLv3+Disable-TLSv1`.

clientcalist {**on** | **off**}
Indicates whether to enable the transmission of a client CA List during the SSL handshake. Transmission of a client CA List is meaningful only when this profile supports a reverse (or server) proxy and when this profile has validation credentials. The default value is `off`.

on Enables the transmission of a client CA List during the SSL handshake.

off Disables the transmission of a client CA List during the SSL handshake.

Guidelines

A cryptographic profile defines a level of SSL service. When you create an SSL proxy profile with the **sslproxy** command, you assign a cryptographic profile to the SSL proxy profile.

Before you create a cryptographic profile to use with an SSL server, use the **certificate** command with the **key** and **idcred** commands to create identification credentials. This set of credentials consists of a certificate, which contains a public key, and the corresponding private key.

A cryptographic profile optionally uses validation credentials to validate certificates that are received from remote SSL peers.

- The SSL client requires validation credentials only when it validates the certificate that is presented by an SSL server. The SSL standard does not require the validation of the server certificate.
- The SSL server requires validation credentials only when it validates certificates that are presented by SSL clients. The SSL standard does not require the validation of SSL clients.

If you want the SSL service to validate received certificates:

1. Use the **valcred** and **certificate** commands to create the validation credentials.
2. Assign the validation credentials to the cryptographic profile.

Assignment of validation credentials to a cryptographic profile mandates that SSL validates the certificate that is presented by the remote peer. If the peer fails to present a certificate on request or presents a certificate that cannot be validated, the cryptographic profile requires the termination of the SSL connection.

Note: In the absence of the `ssl` keyword, the cryptographic profile does no SSL peer authentication.

The **no profile** command deletes only the specified cryptographic profile. The alias names that create the original cryptographic profile are available for use as are the files that contain the certificate and key material that implement the cryptographic profile.

Examples

- Create the `Low` cryptographic profile that uses the `XSSL-1` identification credentials (certificate and private key) to identify the SSL proxy profile. The cryptographic profile specifies no validation of received peer certificates and supports the default cipher list.

```
# profile Low XSSL-1
Creating new crypto profile 'Low'
```

- Create the `Low` cryptographic profile that uses the `XSSL-1` identification credentials to identify the SSL proxy profile. The cryptographic profile specifies no peer validation, supports the default cipher list, and disables SSL Version 2.

```
# profile Low XSSL-1
option-string Disable-SSLv2
Creating new crypto profile 'Low'
```

- Create the `Low` cryptographic profile that uses the `XSSL-1` identification credentials to identify the SSL proxy profile. The cryptographic profile specifies no peer validation, supports the default cipher list, and disables SSL Version 2 and TLS Version 1.0.

```
# profile Low XSSL-1
option-string Disable-SSLv2+Disable-TLSv1
Creating new crypto profile 'Low'
```

- Create the `High` cryptographic profile that uses the `XSSL-2` identification credentials to identify the SSL proxy profile. The cryptographic profile validates

the SSL peer with the TSC-1 validation credentials, and supports symmetric encryption algorithms that include AES and 3DES cipher suites.

```
# profile High XSSL-2
 ssl TSC-1 ciphers HIGH
Creating new crypto profile 'High'
```

- Delete the High cryptographic profile.

```
# no profile High
Crypto Profile 'High' deleted
```

sskey

This command creates an alias for a shared secret key.

Syntax

sskey *alias URL*

no sskey *alias*

Parameters

alias Specifies an alias for the stored shared secret key.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

URL Specifies a local URL that identifies the file that contains the private key.

Guidelines

Use the **sskey** command with the **certificate** and **idcred** commands to create identification credentials that consist of a certificate, which contains a public key, and the corresponding private key.

The **no sskey** command deletes only the alias for the stored shared secret key. The file that contains the actual shared secret key remains on the appliance.

Examples

- Create an alias, alice, for the specified shared secret key.

```
# sskey alice cert:///alicekey
Creating key 'alice'
```

- Delete the alice alias.

```
# no sskey alice
Key 'alice' deleted
```

ssl-server

This command enters SSL Server Profile mode.

Syntax

ssl-servername *name*

no ssl-servername *name*

Parameters

name

Specifies the name of the configuration.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

Guidelines

The **ssl-server** command enters SSL Server Profile mode to create or modify an SSL server profile. An SSL server profile secures connections between clients and the appliance.

To delete an SSL server profile, use the no **ssl-server** command.

test password-map

This command tests the association between an encrypted password alias and a file.

Syntax

test password-map *alias type URL*

Parameters

alias Specifies the name of the candidate alias.

type Identifies the file type. Use the value key or cert.

URL Specifies a local URL that identifies the file that contains the certificate or key.

- If stored in the public cryptographic area, takes the pubcert:*file* form.
- If stored in the private cryptographic area, takes the *file* form.

Guidelines

The **test password-map** command tests the association between an encrypted password alias and a file. Confirms or denies that the alias references the password that protects the file.

Assuming syntactical correctness, testing a key or certificate file that does not require a password succeeds in all cases.

Note: The **test password-map** command cannot be used in a startup configuration. If found, the script ignores the command.

Examples

- Indicates that towson does not reference the encrypted password that protects the dpSupplied.der certificate file.

```
# test password-map towson cert pubcert:dpSupplied.der
Alias 'towson' with file 'pubcert:dpSupplied.der' --> FAIL
```

- Indicates that dundaulk references the encrypted password that protects the dpSupplied.der certificate file.

```
# test password-map dundaulk cert pubcert:dpSupplied.der
Alias 'dundaulk' with file 'pubcert:dpSupplied.der' --> OK
```

- Indicates that columbia does not reference the encrypted password that protects the K2.der key file.

```
# test password-map columbia key K2.der
Alias 'columbia' with file 'K2.der' --> FAIL
```

- Indicates that towson references the encrypted password that protects the K2.der key file.

```
# test password-map towson key K2.der
Alias 'towson' with file 'K2.der' --> OK
```

valcred

This command enters Validation Credentials mode.

Syntax

valcred *name*

no valcred *name*

Parameters

name Specifies the name of the configuration.

The name can contain a maximum of 32 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

Guidelines

The **valcred** command enters Validation Credentials mode. A cryptographic profile optionally uses validation credentials to validate certificates from remote SSL peers.

- Validation credentials are required by an SSL client only when it validates the certificate that is presented by an SSL server. Validation of a server's certificate is not required by the SSL standard.
- Validation Credentials are required by an SSL server only when it validates certificates that are presented by SSL clients. Validation of SSL clients is not required by the SSL standard.

If you want the SSL service to validate received certificates, complete the following procedure.

1. Use the **valcred** command to create the validation credentials.
2. Use the **certificate** command to add certificate alias to the validation credentials.
3. Assign the validation credentials to the cryptographic profile.

The assignment of validation credentials to a cryptographic profile mandates that SSL validates the certificate that is presented by the remote peer. If the peer fails to present a certificate on request or presents a certificate that cannot be validated, the cryptographic profile requires the termination of the SSL connection.

Use the **no valcred** command to delete the validation credentials. The certificate aliases for the validation credentials remain available for use as do as the files that contain the actual certificates.

Date, time, and locale configuration commands

You can use these commands to configure the time, date, and locale on the IBM MQ Appliance.

Included are commands for configuring the appliance to work with an NTP server. To work with date, time, and locale configuration:

1. From the appliance command line, enter global configuration mode:
   ```
   config
   ```
2. From global configuration mode, enter one of the following configuration modes, depending on the feature that you want to configure:
   - For configuring the appliance to work with an NTP server, enter `ntp-service`.
   - For configuring the timezone for the appliance, enter `timezone`.
3. When you have finished, save your configuration:
   ```
   write memory
   ```
4. Type `exit` to leave the configuration mode, then type `exit` again to leave global configuration mode.

custom

This command specifies the name of a custom time zone.

Syntax

custom *name*

Parameters

name Specifies the symbolic name of the custom time zone.

Guidelines

The **custom** command specifies the symbolic name of a custom time zone. This name is appended to local times. The name must be three or more alphabetic characters. If you use any other characters, the time zone becomes Coordinated Universal Time.

daylight-name

This command specifies the name of the time zone when in daylight saving time. This name is appended to the time.

Syntax

daylight-name *name*

Parameters

name Specifies the name of the timezone when in daylight saving time.

Guidelines

The **daylight-name** command specifies the name of the time zone when in daylight saving time. This name is appended to the time. The name must be three or more alphabetic characters. If you use any other characters, the time zone becomes Coordinated Universal Time.

This command is meaningful for custom time zones where daylight saving time rules apply.

daylight-offset

This command sets the offset, in hours, for daylight saving time.

Syntax

daylight-offset *hours*

Parameters

hours Specifies the offset (difference) in hours between daylight saving time and regular time.

Guidelines

The **daylight-offset** command sets the offset, in hours, for daylight saving time. Generally, the difference is 1 hour. A value of 1 means that the clock moves forward or back 1 hour when the time boundary is crossed.

This command is meaningful for custom time zones where daylight saving time rules apply.

daylight-start-day

This command specifies the day of the week when daylight saving time starts.

Syntax

daylight-start-day *day*

Parameters

day Specifies the day of the week when daylight saving time starts. The default value is `Sunday`.

- `Monday`
- `Tuesday`
- `Wednesday`
- `Thursday`
- `Friday`
- `Saturday`
- `Sunday`

Guidelines

The **daylight-start-day** command specifies the day of the week when daylight saving time starts.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time starts, you must also define the time and date with the **daylight-start-hours**, **daylight-start-minutes**, **daylight-start-month**, and **daylight-start-week** commands.

Example

Set the second Sunday in March at 2:30 AM as the start of daylight saving time.

```
# daylight-start-month March
# daylight-start-week 2
# daylight-start-day Sunday
# daylight-start-hours 2
# daylight-start- minutes 30
```

daylight-start-hours

This command specifies the hour in the day when daylight saving time starts.

Syntax

daylight-start-hours *hour*

Parameters

hour Specifies the hour in the day when daylight saving time starts. Enter an value in the range 0 - 23.

Guidelines

The **daylight-start-hours** command specifies the hour in the day when daylight saving time starts. The value uses the 24 hour clock. A setting of 2 is 2 AM; and a setting of 14 is 2 PM.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time starts, you must also define the time and date with the **daylight-start-day**, **daylight-start-minutes**,

daylight-start-month, and **daylight-start-week** commands.

Example

Set the second Sunday in March at 2:30 AM as the start of daylight saving time.

```
# daylight-start-month March
# daylight-start-week 2
# daylight-start-day Sunday
# daylight-start-hours 2
# daylight-start- minutes 30
```

daylight-start-minutes

This command specifies the minute in the hour when daylight saving time starts.

Syntax

daylight-start-minutes *minute*

Parameters

minute

Specifies the minute in the hour when daylight saving time starts. Enter a value in the range 0 - 59.

Guidelines

The **daylight-start-minutes** command specifies the minute in the hour when daylight saving time starts.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time starts, you must also define the time and date with the **daylight-start-hours**, **daylight-start-day**, **daylight-start-month**, and **daylight-start-week** commands.

Example

Set the second Sunday in March at 2:30 AM as the start of daylight saving time.

```
# daylight-start-month March
# daylight-start-week 2
# daylight-start-day Sunday
# daylight-start-hours 2
# daylight-start- minutes 30
```

daylight-start-month

This command specifies the month of the year when daylight saving time starts.

Syntax

daylight-start-month *month*

Parameters

month Specifies the month of the year when daylight saving time starts. The default value is `March`.

- `January`
- `February`

- March
- April
- May
- June
- July
- August
- September
- October
- November
- December

Guidelines

The **daylight-start-month** command specifies the month of the year when daylight saving time starts.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time starts, you must also define the time and date with the **daylight-start-hours**, **daylight-start-minutes**, **daylight-start-day**, and **daylight-start-week** commands.

Example

Set the second Sunday in March at 2:30 AM as the start of daylight saving time.

```
# daylight-start-month March
# daylight-start-week 2
# daylight-start-day Sunday
# daylight-start-hours 2
# daylight-start- minutes 30
```

daylight-start-week

This command specifies the instance of the day in the month when daylight saving time starts.

Syntax

daylight-start-week *instance*

Parameters

instance

Specifies the instance of the day in the month when daylight saving time starts. Enter a value in the range 1 - 5.

Guidelines

The **daylight-start-week** command specifies the instance of the day in the month when daylight saving time starts.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time starts, you must also define the time and date with the **daylight-start-hours**, **daylight-start-minutes**, **daylight-start-day**, and **daylight-start-month** commands.

Example

Set the second Sunday in March at 2:30 AM as the start of daylight saving time.

```
# daylight-start-month March
# daylight-start-week 2
# daylight-start-day Sunday
# daylight-start-hours 2
# daylight-start- minutes 30
```

daylight-stop-day

This command specifies the day of the week when daylight saving time stops.

Syntax

daylight-stop-day *day*

Parameters

day Specifies the day of the week when daylight saving time ends. The default value is Sunday.

- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday
- Sunday

Guidelines

The **daylight-stop-day** command specifies the day of the week when daylight saving time ends.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time ends, you must also define the time and date with the **daylight-stop-hours**, **daylight-stop-minutes**, **daylight-stop-month**, and **daylight-stop-week** commands.

Example

Set the second Sunday in November at 2:30 AM as the end of daylight saving time.

```
# daylight-stop-month November
# daylight-stop-week 2
# daylight-stop-day Sunday
# daylight-stop-hours 2
# daylight-stop- minutes 30
```

daylight-stop-hours

This command specifies the hour in the day when daylight saving time ends.

Syntax

daylight-stop-hours *hour*

Parameters

hour Specifies the hour in the day when daylight saving time ends. Enter a value in the range 0 and 23.

Guidelines

The **daylight-stop-hours** command specifies the hour in the day when daylight saving time ends. The value uses the 24 hour clock. A setting of 2 is 2 AM; and a setting of 14 is 2 PM.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time ends, you must also define the time and date with the **daylight-stop-day**, **daylight-stop-minutes**, **daylight-stop-month**, and **daylight-stop-week** commands.

Example

Set the second Sunday in November at 2:30 AM as the end of daylight saving time.

```
# daylight-stop-month November
# daylight-stop-week 2
# daylight-stop-day Sunday
# daylight-stop-hours 2
# daylight-stop- minutes 30
```

daylight-stop-minutes

This command specifies the minute in the hour when daylight saving time ends.

Syntax

daylight-stop-minutes *minute*

Parameters

minutes
Specifies the minutes of the hour when daylight saving time ends. Enter a value in the range 0 - 59.

Guidelines

The **daylight-stop-minutes** command specifies the minute in the hour when daylight saving time ends.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time ends, you must also define the time and date with the **daylight-stop-hours**, **daylight-stop-day**, **daylight-stop-month**, and **daylight-stop-week** commands.

Example

Set the second Sunday in November at 2:30 AM as the end of daylight saving time.

```
# daylight-stop-month November
# daylight-stop-week 2
# daylight-stop-day Sunday
# daylight-stop-hours 2
# daylight-stop- minutes 30
```

daylight-stop-month

This command specifies the month of the year when daylight saving time ends.

Syntax

daylight-stop-month *month*

Parameters

month Specifies the month of the year when daylight saving time ends. The default value is November.

- January
- February
- March
- April
- May
- June
- July
- August
- September
- October
- November
- December

Guidelines

The **daylight-stop-month** command specifies the month of the year when daylight saving time ends.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time ends, you must also define the time and date with the **daylight-stop-hours**, **daylight-stop-minutes**, **daylight-stop-day**, and **daylight-stop-week** commands.

Example

Set the second Sunday in November at 2:30 AM as the end of daylight saving time.

```
# daylight-stop-month November
# daylight-stop-week 2
# daylight-stop-day Sunday
# daylight-stop-hours 2
# daylight-stop- minutes 30
```

daylight-stop-week

This command specifies the instance of the day in the month when daylight saving time ends.

Syntax

daylight-stop-week *instance*

Parameters

instance

Specifies the instance of the day in the month when daylight saving time ends. Enter a value in the range 1 - 5.

Guidelines

The **daylight-stop-week** command specifies the instance of the day in the month when daylight saving time ends.

This command is meaningful for custom time zones where daylight saving time rules apply.

When you define when daylight saving time ends, you must also define the time and date with the **daylight-stop-hours**, **daylight-stop-minutes**, **daylight-stop-day**, and **daylight-stop-month** commands.

Example

Set the second Sunday in November at 2:30 AM as the end of daylight saving time.

```
# daylight-stop-month November
# daylight-stop-week 2
# daylight-stop-day Sunday
# daylight-stop-hours 2
# daylight-stop- minutes 30
```

direction

This command specifies the direction, relative to Coordinated Universal Time, of the time zone.

Syntax

direction { **East** | **West** }

Parameters

East The direction is east of Coordinated Universal Time. Asia is east.

West The direction is west of Coordinated Universal Time. North America is west. This setting is the default value.

Guidelines

The **direction** command specifies the direction, relative to Coordinated Universal Time, of the time zone.

- When `West`, the offset is added.
- When `East`, the offset is subtracted.

A time zone that is in Coordinated Universal Time can have an offset of 0.

The command is required for custom time zones.

Example

Set an offset of 2 hours and 30 minutes east of Coordinated Universal Time.

```
# direction East
# offset-hours 2
# offset-minutes 30
```

name

This command sets the name of the time zone. This name is appended to the time.

Syntax

name *value*

Parameters

value Specifies the name of the local time zone. The default value is `EST-5EDT`. The appliance provides predefined values and the ability to define a custom time zone.

Value	Description	Meaning
`HST10`	US Hawaii-Aleutian Time	10 hours west of Coordinated Universal Time (UTC), no daylight saving time (DST).
`AKST9AKDT`	US Alaska Time	9 hours west of UTC, United States DST rules.
`PST8PDT`	US Pacific Time	8 hours west of UTC, United States DST rules.
`MST7MDT`	US Mountain Time	7 hours west of UTC, United States DST rules.
`CST6CDT`	US Central Time	6 hours west of UTC, United States DST rules.
`EST5EDT`	US Eastern Time	5 hours west of UTC, United States DST rules.
`AST4ADT`	Atlantic Time	4 hours west of UTC, Canada DST rules.
`UTC`	Coordinated Universal Time	UTC, no DST.
`GMT0BST`	GMT, United Kingdom	UTC, United Kingdom DST rules.
`CET-1CEST`	Central Europe Time	1 hour east of UTC, European Union DST rules.
`EET-2EEST`	Eastern Europe Time	2 hours east of UTC, European Union DST rules.
`MSK-3MSD`	Moscow Time	3 hours east of UTC, Russian DST rules.
`AST-3`	Saudi Arabia	3 hours east of UTC, no DST.
`KRT-5`	Pakistan	5 hours east of UTC, no DST.
`IST-5:30`	India Standard Time	5 hours and 30 minutes east of UTC, no DST.

Value	Description	Meaning
NOVST-6NOVDT	Novosibirsk	6 hours east of UTC, Russian DST rules.
CST-8	China Coast Time	8 hours east of UTC, no DST.
WST-8	Australia Western Time	8 hours east of UTC, no DST.
JST-9	Japan	9 hours east of UTC, no DST.
CST-9:30CDT	Australia Central Time	9 hours and 30 minutes east of UTC. Follows standard Central Australia Daylight savings time rules.
EST-10EDT	Australia Eastern Time	10 hours east of UTC. Follows standard Eastern Australia Daylight savings time rules.
EST-10	Australia Eastern Time, Queensland	10 hours east of UTC, no DST.
Custom	User-defined	Custom time zone that you define with or without DST.

Guidelines

The **name** command sets the name of the time zone. This name is appended to the time.

- Enter a predefined time zone names to set the time zone and its corresponding DST or summer time, values.
- Enter Custom to define a custom time zone with or without DST.

Example

Sets the local time zone to IST-5:30 (India, 5:30 hours east of UTC, no DST).

```
# name IST-5:30
```

offset-hours

This command specifies the offset in hours, relative to Coordinated Universal Time, of the time zone.

Syntax

offset-hours *hours*

Parameters

hours Specifies the offset in hours, relative to Coordinated Universal Time, of the time zone. Enter a value in the range 0 - 12.

Guidelines

The **offset-hours** command specifies the offset in hours, relative to Coordinated Universal Time, of the time zone.

The command is required for custom time zones.

Example

Set an offset of 2 hours and 30 minutes east of Coordinated Universal Time.

```
# direction East
# offset-hours 2
# offset-minutes 30
```

offset-minutes

This command specifies the offset in minutes, relative to Coordinated Universal Time, of the time zone.

Syntax

offset-minutes *minutes*

Parameters

minutes

Specifies the offset in minutes, relative to Coordinated Universal Time, of the time zone. Enter a value in the range 0 - 59.

Guidelines

The **offset-minutes** command specifies the offset in minutes, relative to Coordinated Universal Time, of the time zone.

The command is required for custom time zones.

Example

Set an offset of 2 hours and 30 minutes east of Coordinated Universal Time.

```
# direction East
# offset-hours 2
# offset-minutes 30
```

refresh-interval

This command sets the interval between clock synchronizations.

Syntax

refresh-interval *seconds*

Parameters

seconds

Specifies the number of seconds between clock synchronizations. Enter a value in the range 60 - 86400. The default value is 900.

Guidelines

The **refresh-interval** command sets the interval between clock synchronizations.

Examples

Identify the NTP server and set the clock synchronization interval of 5 minutes.

```
# ntp-service
NTP Service configuration
# remote-server Chronos-1
# refresh-interval 300
```

remote-server

This command identifies an NTP server.

Syntax

remote-server *host*

no remote-server

Parameters

host Identifies the NTP server by host name or IP address.

Guidelines

From the CLI, the appliance supports one NTP server at a time. To designate a new NTP server, use the **no ntp-service** command to delete the current server.

The GUI supports the specification of multiple NTP servers. If you run the **no ntp-service** command, all defined NTP servers are deleted. To delete just one NTP server, use the GUI.

Example

Identify the NTP server and specify a synchronization interval of 5 minutes.

```
# ntp-service
NTP Service configuration
# remote-server Chronos-1
# refresh-interval 300
```

DNS commands

You can use the DNS commands to configure the DNS settings on the IBM MQ Appliance.

The DNS commands can be run from the command line interface in DNS configuration mode. To enter DNS configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:
   `config`
2. From global configuration mode, enter DNS configuration mode:
   `dns`
3. Type `exit` to save your changes and leave DNS configuration mode, then type `exit` again to leave global configuration mode.

ip-preference

This command sets the preferred IP version when the DNS provider publishes both versions of addresses.

Syntax

ip-preference { **4** | **6** }

Parameters

4 Sets the IP version preference to 4. This setting is the default setting.

6 Sets the IP version preference to 6.

Guidelines

The **ip-preference** command sets the preferred IP version that the appliance uses when the DNS provider publishes both supported versions of addresses. If the DNS resolves the host name for a remote destination to both a version 4 address and a version 6 address, this property determines which version of the addresses the appliance uses to attempt the connection to the remote destination. The appliance supports version 4, commonly known as IPv4, and version 6, commonly known as IPv6.

load-balance

This command sets the load-balancing algorithm that the appliance uses to resolve host names.

Syntax

load-balance { **first-alive** | **round-robin** }

Parameters

first-alive
Uses the concept of a primary server and one or more backup servers. When the health state of the primary server is `up`, all connections are forwarded to this server. When the health state of the primary server is `softdown` or `down`, connections are forwarded to back up servers. The primary server is the first server in the members list.

round-robin
Maintains a list of servers and forwards a new connection to the next server on the list. This setting is the default value.

Guidelines

The **load-balance** command sets the load balancing algorithm. For a request to resolve a host name, a server with a health state of `up` is selected from the pool according to the algorithm. The algorithm provides a method to select which server with a health status of `up` receives an incoming client request.

name-server

This command manages local DNS providers.

Syntax

name-server *address*

no name-server *address*

no name-server *

Parameters

address
Specifies the IP address or host name of the DNS server.

***** Indicates all DNS servers.

Guidelines

The **name-server** command manages the list of DNS servers. Use this command to create a list of one or more name servers that the appliance contacts to resolve host names. Each use of the command adds a server to the list.

Note: Unless explicitly instructed, do not change the settings for DNS name servers.

Use the **no name-server** command to delete a DNS provider.

The following syntax is deprecated. For details about this usage, see the online help.

name-server *address* [*UDP*] [*TCP*] [*flags*] [*count*]

This command is equivalent to the Global **ip name-server** command.

Examples

- Add the DNS server at 10.10.10.240 with the default port.

  ```
  # ip name-server 10.10.10.240
  ```

- Delete the specified DNS provider.

  ```
  # no name-server 10.10.10.240
  ```

- Delete all DNS providers.

  ```
  # no name-server *
  ```

retries

This command sets the number of times that the appliance attempts a failed query.

Syntax

retries *attempts*

Parameters

attempts

The number of times that the appliance attempts a failed query. If the query fails, the appliance attempts the query to a different DNS server that is determined by the load balance algorithm. Enter a value in the range 0 - 4294967295. The default value is 2.

Guidelines

The **retries** command sets the number of times that the appliance attempts a failed query before an error is returned. This command is used when the load balancing algorithm is first-alive. Any value set with the **retries** command is ignored when the load balancing algorithm is round-robin.

With the **timeout** command, the **retries** command specifies the number of attempts that the appliance makes to resolve DNS server queries after the initial attempt.

search-domain

This command manages domain-suffixes in the search table for nonqualified domain names.

Syntax

search-domain *domain*

no search-domain *domain*

Parameters

domain

Specifies a base domain name to which a host name can be prefixed.

Guidelines

The **search-domain** command manages domain-suffixes in the search table for nonqualified domain names. Use this command to create a list of one or more domain names that can be added to a host name to resolve host names. Each use of the command adds an entry to the search table.

The appliance attempts to resolve a host name with any domains that are defined with this command. The host name is resolved as soon as a match is found.

Use the **no search-domain** command to delete an entry from the search table.

This command is equivalent to the Global **ip domain** command.

Examples

- Add `datapower.com` to the search table.

  ```
  # search-domain datapower.com
  #
  ```

- Remove `datapower.com` from the search table.

  ```
  # no search-domain datapower.com
  #
  ```

- Add `datapower.com`, `somewhereelse.com`, and `endoftheearth.com` to the search table.

  ```
  # search-domain datapower.com
  # search-domain somewhereelse.com
  # search-domain endoftheearth.com
  # exit
  #
  ```

  The appliance attempts to resolve the host name `loki` as follows:

  1. `loki.datapower.com`
  2. `loki.somewhereelse.com`
  3. `loki.endoftheearth.com`

static-host

This command manages host-address maps.

Syntax

static-host *host address*

no static-host *host*

no static-host *

Parameters

host Specifies the host name.

address
Specifies the IP address of the host.

***** Specifies all hosts.

Guidelines

The **static-host** command manages host-address maps.

Use the **no static-host** command to remove a host-address map or all host-address maps.

This command is equivalent to the Global **ip host** command.

Examples

- Map host `loki-v4` to address `10.10.10.168`.
  ```
  # static-host loki-v4 10.10.10.168
  ```
- Map host `loki-v6` to address `FEDC:BA98:7654:3210:C:BA98:7654:3210`.
  ```
  # static-host loki-v6 FEDC:BA98:7654:3210:C:BA98:7654:3210
  ```
- Delete the map for host `loki-v4`.
  ```
  # no static-host loki-v4
  ```
- Delete all maps.
  ```
  # no static-host *
  ```

timeout

This command sets the time to wait before the next query attempt.

Syntax

timeout *seconds*

Parameters

seconds
Specifies the number of seconds to wait for a response from a remote DNS server. If the query fails, the appliance attempts the query to a different DNS server that is determined by the load balance algorithm. The default value is 5.

Guidelines

The **timeout** command sets the time to wait before the next query attempt. Use this command only when the load-balancing algorithm is first-alive. Any value set with the **timeout** command is ignored when the load balancing algorithm is round-robin.

With the **retries** command, the **timeout** command specifies the amount of time that the appliance attempts to resolve DNS server queries.

Ethernet commands

You can use the Ethernet commands to configure the Ethernet interfaces on the IBM MQ Appliance.

The Ethernet commands can be run from the command line interface in Ethernet configuration mode. To enter Ethernet configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:
   ```
   config
   ```
2. From global configuration mode, enter Ethernet configuration mode:
   ```
   ethernet name
   ```
   where *name* is the name of the Ethernet port that you want to configure.
3. Type `exit` to leave the configuration mode and save your changes, then type `exit` again to leave global configuration mode.

You cannot aggregate Ethernet links that are used in high availability or disaster recovery configurations.

If you change the IP addresses of Ethernet links that are used in high availability or disaster recovery configurations, the configuration will cease to function. See "Changing IP addresses in high availability configurations" on page 129 and "Changing IP addresses in disaster recovery configurations" on page 146

.

flow-control

This command sets the flow control mode of the Ethernet interface.

Syntax

flow-control { **auto** | **disabled** }

flow-control { **auto** | **disabled** | **tx** | **rx** | **full** }

Parameters

auto For interfaces that support autonegotiation, performs standard IEEE 802.3 autonegotiation for flow control. This setting is the default value.

disabled
Disables flow control. The interface does not send flow control `PAUSE` frames and ignores received `PAUSE` frames.

Note: Do not select this option unless you are working with IBM Support.

tx Enables transmit mode. The interface transmits flow control PAUSE frames but ignores the received PAUSE frames.

rx Enables receive mode. The interface honors received flow control PAUSE frames but does not transmit PAUSE frames.

full Enables full flow control. The interface transmits flow control PAUSE frames and honors the received PAUSE frames.

Guidelines

The **flow-control** command specifies the flow control mode of the Ethernet interface. Disable flow control only at the direction of IBM Support.

Example

Set the flow control mode of Ethernet interface eth20 to disabled.

```
# interface eth20
Interface configuration mode (eth20)
(config-if[eth20]) # flow-control disabled
```

disable-ethernet-hardware-offload

This command manages the temporary disabling of hardware offload.

Syntax

disable-ethernet-hardware-offload

Guidelines

The **disable-ethernet-hardware-offload** command manages the temporary disabling of hardware offload. This command does not modify the interface configuration. Use the **disable-ethernet-hardware-offload** command to disable hardware offload temporarily. In rare situations, hardware offload can cause stability issues.

The hardware offload function, as set with the **hardware-offload** command, controls whether to offload TCP/IP packet processing of Ethernet device drivers and chips.

Example

Temporarily disable hardware offload.

```
# disable-ethernet-hardware-offload
```

ip-address

This command assigns the primary network address for the Ethernet interface.

Syntax

ip-address *address*

Parameters

address

Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

Guidelines

The **ip-address** command assigns the primary network address to the interface. The network address is an IP address with its subnet mask.

To assign secondary, or auxiliary, IP addresses, use the **ip-secondary-address** command.

This command is meaningful except in the following situations:

- You use the **link-aggregation-mode** command to make the interface part of an aggregate interface.
- You use the **ip-config-mode** command to identify autoconfiguration with DHCP or SLAAC.

Examples

- Assign an IP address in version 4 format.

  ```
  # ip-address 192.168.7.6/27
  ```

- Assign an IP address in version 6 format.

  ```
  # ip-address 2001:0db8:3c4d:0015::abcd:ef12/34
  ```

ip-config-mode

This command identifies the configuration mode for the Ethernet interface.

Syntax

ip-config-mode { **static** | **dhcp** | **slaac** }

Parameters

static Indicates a static, manual configuration. This setting is the default value.

dhcp Indicates IPv4 autoconfiguration with DHCP.

slaac Indicates IPv6 autoconfiguration with SLAAC.

Guidelines

The **ip-config-mode** command identifies the configuration mode of the interface.

- With the `static` keyword, define the configuration for the interface as provided by your network administrator.
  - Use the **ip-address** command to assign the primary network address.
  - Use the **ip-secondary-address** command to manage secondary, or auxiliary, network addresses.
  - Use the **ipv4-default-gateway** command to assign the default IPv4 gateway.
  - Use the **ipv6-default-gateway** command to assign the default IPv6 gateway.
  - Use the **ip-route** command to manage static routes in the routing table.
- With the `dhcp` keyword, the appliance ignores configuration data about the physical interface.
- With the `slaac` keyword, the appliance ignores configuration data about the physical interface.

This command is meaningful only when you do not use the **link-aggregation-mode** command to make interface part of an aggregate interface.

Examples

- Change the configuration mode to IPv4 autoconfiguration with DHCP.

  ```
  # ip-config-mode dhcp
  ```

- Change the configuration mode to manual configuration.

  ```
  # ip-config-mode static
  ```

ip-route

This command manages static routes in the routing table for the Ethernet interface.

Syntax

Add a static route

ip-route *address next-hop-address* [*metric*]

Delete a static route

no ip-route *address next-hop-address*

Parameters

address

Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

next-hop-address

Specifies the IP address of the next-hop router.

metric Optionally specifies the preference for the route. The lesser the value, the more preferred the route. For each IP family, the supported range differs.

- For IPv4, enter a value in the range 0 - 255. The default value is 0.
- For IPv6, enter a value in the range 0 - 65536. The default value is 512.

Guidelines

The **ip-route** command manages static routes in the routing table. Issue this command for each static route to add to the routing table.

To delete a static route, use the **no ip-route** command. Issue this command for each static route to delete from the routing table.

This command is meaningful except in the following situations:

- You use the **link-aggregation-mode** command to make the interface part of an aggregate interface.
- You use the **ip-config-mode** command to identify autoconfiguration with DHCP or SLAAC.

Examples

- Add a static route to the routing table (subnet 10.10.10.224 via next-hop router 192.168.1.100). The metric for the route is 0, the default value for IPv4, which is the most preferred route.

  ```
  # ip-route 10.10.10.0/27 192.168.1.100
  ```

- Delete a static route from the routing table (subnet 10.10.10.224 via next-hop router 192.168.1.100).

  ```
  # no ip-route 10.10.10.0/27 192.168.1.100
  ```

ip-secondary-address

This command manages secondary network addresses for the Ethernet interface.

Syntax

Add a secondary address

ip-secondary-address *address*

Remove a secondary address

no ip-secondary-address *address*

Remove all secondary addresses

no ip-secondary-address

Parameters

address

Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

Guidelines

The **ip-secondary-address** command manages secondary network addresses for the current interface. The network address is the IP address and its subnet mask. A secondary IP address is a bind address. The secondary IP address is used only as a source IP address when it responds to incoming traffic to the secondary IP address.

To create the primary IP address, use the **ip-address** command.

To remove secondary IP addresses, use the **no ip-secondary-address** command.

This command is meaningful except in the following situations:

- You use the **link-aggregation-mode** command to make the interface part of an aggregate interface.
- You use the **ip-config-mode** command to identify autoconfiguration with DHCP or SLAAC.

Examples

- Add 192.168.7.6/27 as a secondary IP address to the interface.

  ```
  # ip-secondary-address 192.168.7.6/27
  ```

- Remove 192.168.7.6/27 as a secondary IP address.

  ```
  # no ip-secondary-address 192.168.7.6/27
  ```

- Remove all secondary IP addresses.

  ```
  # no ip-secondary-address
  ```

ipv4-default-gateway

This command designates the default IPv4 gateway for the Ethernet interface.

Syntax

Designates the default IPv4 gateway
ipv4-default-gateway *address*

Deletes the default IPv4 gateway
no ipv4-default-gateway

Parameters

address
Specifies the IP address of the default IPv4 gateway.

Guidelines

The **ipv4-default-gateway** command designates the default IPv4 gateway that the interface can reach. If the interface supports both IP families, use the **ipv6-default-gateway** command to designate the default IPv6 gateway.

Use the **no ipv4-default-gateway** command to delete the default IPv4 gateway.

This command is meaningful except in the following situations:

- You use the **link-aggregation-mode** command to make the interface part of an aggregate interface.
- You use the **ip-config-mode** command to identify autoconfiguration with DHCP or SLAAC.

ipv6-dadtransmits

This command sets the number of IPv6 duplication address detection attempts for the Ethernet interface.

Syntax

ipv6-dadtransmits *attempts*

Parameters

attempts
Specifies the number of attempts. The default value is 1.

Guidelines

The **ipv6-dadtransmits** command sets the number of IPv6 duplication address detection (DAD) attempts. This command is relevant for only IPv6 addresses on the appliance.

If you specify more than one attempt, use the **ipv6-nd-retransmit-timer** command to set the interval between attempts.

This command is meaningful except when you use the **link-aggregation-mode** command to make the interface part of an aggregate interface.

ipv6-default-gateway

This command designates the default IPv6 gateway for the Ethernet interface.

Syntax

Designate the default IPv6 gateway
ipv6-default-gateway *address*

Delete the default IPv6 gateway
no ipv6-default-gateway

Parameters

address
Specifies the IP address of the default IPv6 gateway.

Guidelines

The **ipv6-default-gateway** command designates the default IPv6 gateway that the interface can reach. Define a default IPv6 gateway if you defined IPv6 IP addresses.

If the interface supports both IP families, use the **ipv4-default-gateway** command to designate the default IPv4 gateway.

Use the **no ipv6-default-gateway** command to delete the default IPv6 gateway.

This command is meaningful except in the following situations:

- You use the **link-aggregation-mode** command to make the interface part of an aggregate interface.
- You use the **ip-config-mode** command to identify autoconfiguration with DHCP or SLAAC.

ipv6-nd-retransmit-timer

This command sets the interval between IPv6 neighbor discovery attempts for the Ethernet interface.

Syntax

ipv6-nd-retransmit-timer *milliseconds*

Parameters

milliseconds
Specifies the interval between attempts in milliseconds. The default value is 1000.

Guidelines

The **ipv6-nd-retransmit-timer** command sets the interval neighbor discovery attempts. This command is relevant for only when the interface uses IPv6 addresses.

This command is meaningful except when you use the **link-aggregation-mode** command to make the interface part of an aggregate interface.

link

This command manages Ethernet interfaces in the aggregate interface.

Syntax

Add an Ethernet interface
link *name*

Delete an Ethernet interface
no link *name*

Parameters

name Specifies the name of an Ethernet interface.

Guidelines

The **link** command manages Ethernet interfaces in the aggregate interface.

- For each Ethernet interface to add to the aggregate interface, use the **link** command.
- For each Ethernet interface to delete from the aggregate interface, use the **no link** command.

You cannot add members to an aggregate interface in the following situations:

- The Ethernet interface is part of a VLAN interface
- The Ethernet interface is part of another aggregate interface
- The Ethernet interface is not enabled for link aggregation with the **link-aggregate-mode** command in Ethernet mode.

Examples

Add the eth10, eth11, and eth12 Ethernet interfaces to the aggregate interface.

```
# link eth10
# link eth11
# link eth12
```

Deletes the eth12 Ethernet interface from the aggregate interface.

```
# no link eth12
```

mac-address

This command changes the MAC address for the Ethernet interface.

Syntax

mac-address *address*

Parameters

address
Specifies the 48-bit MAC address in hex.

Guidelines

The **mac-address** command changes the MAC address of the Ethernet interface. By default, the appliance uses “burned-in” MAC addresses from the network interface controller.

Example

Change the "burned-in" MAC address to a nondefault value.

```
# mac-address 00:11:22:aa:bb:cc
```

mode

This command sets the interface speed and direction.

Syntax

mode { **Auto** | **10baseT-FD** | **10baseT-HD** | **100baseTx-FD** | **100baseTx-HD** | **1000baseTx-FD** | **10000baseTx-FD** }

Parameters

Auto For interfaces that do autonegotiation, the appliance uses standard IEEE 802.3 autonegotiation for interface speed and direction. Preference is given to the highest speed. Preference is for full-duplex over half-duplex. This setting is the default value.

10baseT-FD
Advertises 10BASE-T PHY (10 Mbps) in full-duplex mode.

10baseT-HD
Advertises 10BASE-T PHY (10 Mbps) in half-duplex mode.

100baseTx-FD
Advertises 100BASE-TX PHY (100 Mbps) in full-duplex mode.

100baseTx-HD
Advertises 100BASE-TX PHY (100 Mbps) in half-duplex mode.

1000baseTx-FD
Advertises 1000BASE-T PHY (1 Gbps) in full-duplex mode.

10000baseTx-FD
Advertises 10000BASE-T PHY (10 Gbps) in full-duplex mode.

Guidelines

The **mode** command specifies the operational mode for the interface. The operational mode is the interface speed and direction. Generally you can retain the default value.

On some appliances, you cannot use this command to modify the operational mode. Because some network equipment might not negotiate properly, you can use this command to set speed and direction explicitly. If you manually configure one end of the link, you must manually configure the other end of the link to the same setting.

On physical appliances, you cannot modify the operational mode for Ethernet interfaces that use 10-gigabit ports. On blades, you cannot modify the operational mode on any Ethernet interface.

mtu

This command sets the maximum transmission unit of the Ethernet interface.

Syntax

mtu *bytes*

Parameters

bytes Specifies the maximum size in bytes. Enter a value in the range 576 - 16128. The default value is 1500.

Guidelines

The **mtu** command sets the maximum transmission unit (MTU) for the Ethernet interface. The MTU is determined regardless of the length of the layer 2 encapsulation.

- When the Ethernet interface is part of a VLAN interface, the MTU of the VLAN interface cannot be greater than the MTU for the Ethernet interface.
- When the Ethernet interface is part of an aggregate interface, the MTU of the aggregate interface overrides the MTU of the Ethernet interface.

Example

Set the MTU to 4 KB.

```
# mtu 4096
```

packet-capture

This command manages a packet-capture for the Ethernet interface session.

Syntax

Start a packet-capture session

packet-capture *file seconds KB* ["*expression*"]

Stop a packet-capture session

no packet-capture *file*

Parameters

file Specifies the file name for the packet capture. You can simultaneously capture packets on multiple interfaces by specifying a different file name for each interface.

seconds

Specifies the maximum duration of the packet-capture session in seconds. Enter a value in the range 5 - 86400. The special value of **-**1 indicates that the packet capture is continuous and completes when it reaches the maximum file size or until you issue the **no packet-capture** command.

KB Specifies the maximum size of the file in KB. Enter a value in the range 10 - 500000.

expression

Optionally specifies the expression that filters the packet capture. Enclose the expression in double quotation marks.

Guidelines

The **packet-capture** command manages a packet-capture session on the current interface. The data from the session is saved in the `pcap` format. To interpret the packet, use a network protocol analyzer.

Use the **no packet-capture** command to stop a packet-capture session.

Examples

- Start a timed packet-capture session that writes data to the temporary:/// capture-1 file. The session completes either after 30 minutes or when the file contains 2500 KB, whichever occurs first.

```
# packet-capture temporary:///capture-1 1800 2500
Trace begun.
#
```

- Start a timed packet-capture session that writes data to the temporary:/// capture-2 file. The session records only packets where 53 is the destination port. The session completes either after 30 minutes or when the file contains 2500 KB, whichever occurs first.

```
# packet-capture temporary:///capture-2 1800 2500 "dst port 53"
Trace begun.
#
```

- Start a continuous packet-capture session that writes data to the temporary:///capture-3 file. The session completes either when it contains 50000 KB or when you stop it.

```
# packet-capture temporary:///capture-3 -1 50000
Trace begun.
#
```

- Stop the packet-capture session that writes data to the temporary:///capture-3 file.

```
# packet-capture temporary:///capture-3
Continuous packet capture to temporary:///capture-3 on interface stopped.
#
```

File commands

You can use the file commands to move files on and off the IBM MQ Appliance and list the contents of directories.

The file commands can be run from the command line interface in configuration mode. To enter configuration mode, type config.

copy

This command copies a file to or from the appliance.

Syntax

copy [**-f**] *source destination* [*manager*]

Parameters

-f Overwrites an existing file when one with the same name exists. When omitted, an attempt to save a file with the same name as an existing file results in a confirmation prompt.

source
destination

Specifies the locations as a URL that identifies the source file and target destination.

- When the source or destination is the appliance, use the *directory*:///*file* format.
- If the source file or target destination is remote and the transport protocol is SCP or SFTP, use a format that is RFC 1738 compliant.

To use an absolute path.

`scp://user@host:port//file_path`

`sftp://user@host:port//file_path`

To use a path that is relative to the user's home directory.

`scp://user@host:port/file_path`

`sftp://user@host:port/file_path`

Where:

host Specifies the fully qualified host name or IP address of the remote server. If DNS is enabled, the host name.

port Specifies the listening port on the remote server.

After you enter the command, the appliance prompts you for the remote login password.

manager

Specifies the name of an XML manager that defines the user agent for basic authentication. When omitted, the appliance uses the basic authentication settings for the `default` XML manager.

Guidelines

The **copy** command transfers files to or from the DataPower® appliance. You must issue this command from the appliance.

- The optional `-f` parameter forces an unconditional copy. When provided, the command does not warn of possible file overwrites.
- The optional *manager* parameter defines the basic authentication configuration to use. When provided, the command uses the user agent for this XML manager instead of the `default` XML manager.

When the source file or target destination is remote to the appliance, this command supports only the following protocols:

- HTTP
- HTTPS
- Secure Copy (SCP)
- Secured File Transfer Protocol (SFTP)

Note: If you use SCP or SFTP to copy multiple files at the same time, a system error might occur. The appliance supports only one SCP or SFTP connection at a time. Issue the command again if you encounter a system error.

To send a file from the appliance as an email, use the global **send file** command.

Restriction: When you use the **copy** command, be aware of the following restrictions.

- You cannot copy files from the `cert:` directory.
- You cannot copy files to the `audit:`, `logstore:`, or `logtemp:` directory.

Examples

- Use HTTP to copy a file from the specified location to the `image:` directory.

```
(config)# copy http://host/image.crypt image:///image.crypt
File copy success (1534897 bytes transferred)
```

- Use HTTP over SSL to copy a file from the specified location to the `image:` directory.

```
(config)# copy https://host/image.crypt image:///image.crypt
File copy success (1534897 bytes transferred)
```

- Use HTTP to copy a file from the specified location to the `local:` directory with the basic authentication credentials in the `sec` XML manager.

```
(config)# copy http://host/sec/stock.wsdl local:///stock.wsdl sec
File copy success (2022 bytes copied)
```

- Use SCP to copy a file from the specified location to the `store:` directory.

```
(config)# copy scp://jrb@10.10.1.159//XML/stylesheets/InitialConvert.xsl
store:///InitialConvert.xsl
     Password: yetanotherpassword
File copy success
```

- Use SCP to copy a file from the `logstore:` directory to the specified remote target (identified by a qualified host name).

```
(config)# copy logstore:///Week1.log scp://jrb@ragnarok.datapower.com//LOGS/Week1.log
     Password: yetanotherpassword
File copy success
```

- Use SFTP to copy a file from the specified location to the `store:` directory.

```
(config)# copy sftp://jrb@10.10.1.159//XML/stylesheets/InitalConvert.xsl
 store:///InitalConvert.xsl
     Password: yetanotherpassword
File copy success
```

- Use SFTP to copy a file from the `logstore:` directory to the specified remote target.

```
(config)# copy logstore:///Week1.log sftp://jrb@10.10.1.159//LOGS/x/Week1.log
     Password: yetanotherpassword
File copy success
```

- Copy the `startup-config` file from the `config:` directory to the `local:` directory.

```
(config)# copy config:///startup-config local:///startup-config
file copy successful (2347 bytes transferred)
```

delete

This command deletes a file from the appliance.

Syntax

delete *URL*

Parameters

URL

Specifies the location as a URL of the file to delete in the *directory*`:///`*file* format.

Guidelines

The **delete** command deletes a file on the DataPower appliance. The deletion of a file is permanent. After you delete a file, it cannot be recovered.

Attention: The **delete** command does not prompt for confirmation.

Examples

- Delete the `startup-config-deprecated` file from the `store:` directory.

```
(config)# delete store:///startup-config-deprecated
```

- Delete the `betaImage` file from the `image:` directory.

```
(config)# delete image:///betaImage
```

dir

This command lists the contents of a directory.

Syntax

dir *directory*

Parameters

directory
Specifies a directory on the appliance.

Examples

- List the contents of the `config:` directory.

```
(config)# dir config:

File Name         Last Modified              Size
---------         -------------              ----
unicenter.cfg     Mon Jul  9 11:09:36 2007   3411
autoconfig.cfg    Mon Jul  9 14:20:27 2007   20907

89.2 MB available to config:
```

- List the contents of the `msgcat` subdirectory of the `store:` directory.

```
(config)# dir store:\\\msgcat

File Name         Last Modified              Size
---------         -------------              ----
crypto.xml        Mon Jul  9 11:09:26 2007   179069
dplane.xml        Mon Jul  9 11:09:26 2007   299644
...
xslt.xml          Mon Jul  9 11:09:26 2007   10233

89.2 MB available to store:\msgcat
```

move

This command moves a file from one directory to another.

Syntax

move [**-f**] *source* destination

Parameters

-f Overwrites an existing file when one with the same name exists. When omitted, an attempt to save a file with the same name as an existing file results in a confirmation prompt.

source
destination
Specifies the locations as a URL that identifies the source file and target destination in the *directory*:///*file* format.

Guidelines

You can use the **move** command to transfer a file to or from a directory.

Restriction: You cannot use the **move** command to copy a file from the private cryptographic area, such as the `cert:` directory.

Examples

- Move a file from the `config:` directory to the `store:` directory.

  ```
  (config)# move config:///startup-config store:///archiveConfig-10
  ```

- Rename a file.

  ```
  (config)# move config:///startup-config config:///archiveConfig-10
  ```

Global configuration commands

You can use the global configuration commands to manage configuring the IBM MQ Appliance.

You enter the commands in configuration mode. Type `config` to enter configuration mode.

appliance-quiesce

This command changes the operational state of the domains, services, and handlers on the appliance to the `down` state in a controlled manner.

Syntax

appliance-quiesce *seconds*

Parameters

seconds

Specifies the time to wait for all transactions to complete. Specify a long enough timeout value to handle long living transactions. The minimum value is 60.

Guidelines

The **appliance-quiesce** command changes all domains, except the `default` domain, to the `down` operational state. The command changes all services and all handlers to the `down` operational state. A quiesce waits for at most the specified time for transactions to complete and proceeds to notify quiesce complete.

This command does not quiesce management interfaces, such as the Telnet service, the web management service, and the XML management interface. The Sysplex Distributor Target Control Service is considered a management interface.

Example

Quiesce all domains, services, and handlers except the `default` domain. Waits for at most 2 minutes for transactions to complete.

```
# appliance-quiesce 120
#
```

appliance-unquiesce

This command changes quiesced domains, services, and handlers on the appliance to the `up` operational state.

Syntax

appliance-unquiesce

Guidelines

The **appliance-unquiesce** command changes quiesced domains, services, and handlers to the `up` operational state. The running configuration reflects any changes that were applied when the appliance was in the quiesced state.

Unquiesce must be done at the same level as the quiesce. For example, you cannot unquiesce a domain if the appliance is quiesced. If you create a domain, service, or handler while the appliance is quiesced, it does not change to the `up` operational state until the appliance is unquiesced.

Examples

Unquiesce the quiesced appliance.

```
# appliance-unquiesce
#
```

audit-log-settings

This command enters Audit Log Settings mode.

Syntax

Enter Audit Log Settings mode

audit-log-settings

Guidelines

The **audit-log-settings** command enters Audit Log Settings mode. In this mode, you can set the size and the number of rotations of the audit log.

backup

This command generates an archive file of the configuration data for specific domains.

Syntax

backup *file* [*domain*]

Parameters

file Specifies the name of the archive file.

domain

Specifies the target domains to back up. In the absence of this argument, all domains are backed up. This argument is available only from the `default` domain.

Guidelines

The file name does not need to include the `.zip` extension. Without the extension, the firmware appends it to the provided name.

The behavior and syntax of the **backup** command vary depending on whether the command is entered from within the default domain or an application domain.

- From the `default` domain, the archive file is in the `export:` directory.
- From an application domain, the archive file is contained in a subdirectory that is named after the domain. For example, from the `alphaDelivery` domain, the archive file is in the `export:` directory of the `alphaDelivery` domain.

The domain name property is available in the `default` domain only. If the argument consists of a list of target domains, use a single space to separate entries.

The **backup** command cannot be used to filter specific objects and domain resources from the archive file. This feature is available with the **Export Objects and Files** option from the **Export** utility.

The **backup** command is equivalent to the **Create a backup of the entire system** or **Create a backup of one or more application domain** options of the **Export** utility.

Note: The following objects are never exported:

- User accounts
- Certificates
- Keys

The following files are never exported:

- Log files
- Firmware files

To ensure that all other objects and files are exported, use the `admin` account or the `dp-admin` account on XI50z. For any other user, only objects and files that are accessible to that user are included in the export package.

Examples

- From the `default` domain, back up all domains to `export:///all.zip`. Use the `dir export:` command to confirm file creation.

```
# backup all.zip
Backup to all.zip scheduled (may take a few minutes to complete)
```

- From the `default` domain, back up the `default` domain to `export:///default.zip`. Use the `dir export:` command to confirm file creation.

```
# backup default default
Backup to default.zip scheduled (may take a few minutes to complete)
```

- From the `default` domain, back up the `wsProxyDev` and `wsdmDev` domains to `export:///project.zip`. Use the `dir export:` command to confirm file creation.

```
# backup project wsProxyDev wsdmDev
Backup to project.zip scheduled (may take a few minutes to complete)
```

- From the `wsProxyDev` domain, back up that domain to `export:///SIP-1.zip`. To confirm file creation, use the `dir export:` command from in the `wsProxyDev` domain, or use the `dir export:wsProxyDev` command from in the `default` domain.

```
# switch domain wsProxyDev
# backup SIP-1
Backup to SIP-1.zip scheduled (may take a few minutes to complete)
```

clear dns-cache

This command clears the DNS cache.

Syntax

clear dns-cache

Examples

Clear the DNS cache.

```
# clear dns-cache
Cleared DNS cache
#
```

clear rbm cache

This command clears all cached RBM authentication data.

Syntax

clear rbm cache

Examples

Clear cached RBM authentication data.

```
# clear rbm cache
Cleared RBM cache
#
```

dns

This command enters DNS Settings mode.

Syntax

dns

no dns

Context

On XI50z, the DNS settings are read-only.

Guidelines

The **dns** command enters DNS Settings mode to define DNS settings.

Use the **no dns** command to disable DNS services.

Examples

- Enter DNS Settings mode.

  ```
  # dns
  DNS Settings mode
  ```

- Disable DNS.

  ```
  # no dns
  ```

ethernet

This command enters Ethernet Interface mode to manage the configuration of Ethernet interfaces.

Before you configure an interface, obtain the essential network data from your network team.

Syntax

Accesses the configuration of an Ethernet interface

ethernet *name*

Deletes the configuration for an Ethernet interface

no ethernet *name*

Parameters

name Specifies the name of the configuration. The **show link** command lists the supported Ethernet interfaces.

Guidelines

The **ethernet** command enters Ethernet mode for a specific Ethernet interface.

If you use a name for an Ethernet interface that your appliance type does not support, you create that configuration. However, this configuration is not associated with the appliance and cannot be used for traffic. When you import a configuration from another appliance type, the import operation can create these unsupported Ethernet interface configurations.

- If you accidentally create a configuration, use the **no ethernet** command to delete it.
- When the configuration is created from an import operation, you might need the configuration details.
  - If you need these details, modify the configuration for supported Ethernet interfaces. After you modify the configurations, use the **no ethernet** command to delete the configuration for unsupported Ethernet interfaces.
  - If you do not need these details, use the **no ethernet** command to delete the configuration for unsupported Ethernet interfaces.

The **no ethernet** command deletes the configuration for an Ethernet interface. You can delete the configuration for supported interfaces. If you delete the configuration for a supported interface, you can use a following approach to have a supported interface that can be used for traffic.

- You can create the configuration again.
- Restart the appliance to re-recreate the interface.

To disable an Ethernet interface, use the **admin-state** command in Ethernet interface mode.

Examples

- Enter Interface mode for the `eth10` Ethernet interface.

```
# ethernet eth10
Modify Ethernet Interface configuration
```

- Disable the `eth10` Ethernet interface.

```
# ethernet eth10
Modify Ethernet Interface configuration
# admin-state disable
# exit
```

- Delete the configuration for the `eth5` Ethernet interface.

```
# no ethernet eth5
```

host-alias

This command enters Host Alias mode to map an IP address to an alias.

Syntax

host-alias *alias*

no host-alias *alias*

Parameters

alias Specifies the alias to assign to the specified IP address.

Guidelines

Use the **no host-alias** command to remove an alias map.

ipmi-lan-channel

This command enters IPMI LAN Channel mode.

Syntax

Defines the LAN channel.

```
ipmi-lan-channel mgt0
```

Deletes the LAN channel.

```
no ipmi-lan-channel mgt0
```

Availability

Physical appliances only.

Guidelines

The **ipmi-lan-channel** command enters IPMI LAN Channel mode for the `mgt0` interface. While in the mode, define the LAN channel. The Intelligent Platform Management Interface (IPMI) LAN channel must be on the `mgt0` interface, which is accessible over only the `mgt0` physical interface on the appliance.

An IPMI LAN channel allows access to the Baseboard Management Controller (BMC) on the appliance over a LAN. IPMI allows remote management access and can provide serial over LAN to the console serial port.

CAUTION:
If you enable serial over LAN support for the IPMI LAN channel, an IPMI user who connects through an IPMI 2.0 client has access to the serial console port.

- **An IPMI user can connect through an IPMI 2.0 client independent of the state of the appliance. The only time that an IPMI user cannot connect to the appliance is when the appliance is disconnected from AC power.**
- **An IPMI user has higher priority than a user who is directly connected to the serial port.**
- **If a user is directly connected to the serial port, the IPMI user usurps the current serial session and does not need to log in to the appliance. If no user is directly connected to the serial port, the IPMI user must log in to the appliance.**
- **There can be only one serial user. The IPMI user suspends the serial session of the user who is directly connected to the serial port until the remote IPMI user closes (deactivates) the serial session. When the IPMI user closes the serial session, the session for the user who is directly connected to the serial port resumes.**
- **If an IPMI user is connected to the appliance and you need to use the serial port, you must unplug the Ethernet cable from the `MGT0` port. After you unplug the cable, wait for the serial port to become available, which can take up to 20 minutes.**

Use the **`no ipmi-lan-channel`** command to delete the IPMI LAN channel.

`ipmi-user`

This command enters IPMI User mode.

Syntax

Creates or modifies an IPMI user.

`ipmi-user` *name*

Deletes an IPMI user.

`no ipmi-user` *name*

Availability

Physical appliances only.

Parameters

name Specifies the name of the IPMI user. The name must be 16 characters or less.

Guidelines

The **`ipmi-user`** command enters IPMI User mode. While in the mode, define an Intelligent Platform Management Interface (IPMI) user.

An IPMI user can create, change, or destroy user authentication records in the Baseboard Management Controller (BMC). Authentication records allow users to communicate with IPMI protocols over external channels, such as an IPMI LAN channel. On the IBM MQ Appliance, there can be eight IPMI users.

Note: If you apply the configuration but do not save it, the IPMI user data is written to the BMC but not saved to the startup configuration. In this situation, the IPMI user can connect to the serial port over the IPMI LAN channel.

Use the **no ipmi-user** command to remove an IPMI user.

known-host

This command adds or removes an SSH peer as an SSH known host.

Syntax

known-host *host* **ssh-rsa** *key*

no known-host *host*

Parameters

host Specifies the fully-qualified host name or IP address for the peer. For example:

```
ragnarok.datapower.com
10.97.111.108
```

ssh-rsa
Identifies RSA as the key type.

key Specifies the host public key for the peer. For example:

```
AAAAB3NzaC1yc2EAAAABIwAAAIEA1J/99rRvdZmVvkaKvcG2a+PeCm25
p8OJl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfRl
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi0lud5Pl7+BXQ
VpPbyVujoHINCrx0k/z7Qpkozb4qZd8==
```

Guidelines

The **known-host** command adds an SSH peer as an SSH known host.

The **no known-host** command removes an SSH peer as an SSH known host.

Examples

- Add `ragnarok.datapower.com` by host name as an SSH known host.

```
# known-host ragnarok.datapower.com ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEA1J/99rRvdZmVvkaKvcG2a+PeCm25
p8OJl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfRl
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi0lud5Pl7+BXQ
VpPbyVujoHINCrx0k/z7Qpkozb4qZd8==
#
```

- Add `ragnarok.datapower.com` by IP address as an SSH known host.

```
# known-host 10.97.111.108 ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEA1J/99rRvdZmVvkaKvcG2a+PeCm25
p8OJl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfRl
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi0lud5Pl7+BXQ
VpPbyVujoHINCrx0k/z7Qpkozb4qZd8==
#
```

- Remove `ragnarok.datapower.com` by IP address as an SSH known host.

```
# no known-host 10.97.111.108
#
```

language

This command enters the Language mode to set the administrative state of a language.

Syntax

language *locale*

Parameters

locale Specifies the operating language.

de German.

en English.

es Spanish.

fr French.

it Italian.

ja Japanese.

ko Korean.

pt_BR Brazilian Portuguese.

ru Russian.

zh_CN Simplified Chinese.

zh_TW Chinese (Taiwan).

Guidelines

The **language** command enters Language mode for the specified language.

While in this mode, use the **admin-state** command to enable or disable the language. The **admin-state** command is the only command available in this mode.

- When a language has an administrative state of `enabled`, its operational state is `up`.
- When a language has an administrative state of `disabled`, its operational state is `down`.

Language enablement is for the following purposes:

- Control the operating language of the appliance that you can set with the **locale** command in System Settings mode.
- Support native languages for browsers that connect to the appliance with a web management interface.

Examples

- Enable a locale.

```
# language fr
Modify Language configuration
# admin-state enabled
```

- Disable a locale.

```
# language de
Modify Language configuration
# admin-state disabled
```

link-aggregation

This command enters Link Aggregation mode to manage the configuration of aggregate interfaces.

Before you configure an interface, obtain the essential network data from your network team.

Syntax

Enters the mode to create or modify an aggregate interface configuration
link-aggregation *name*

Deletes an aggregate interface configuration
no link-aggregation *name*

Parameters

name Specifies the name for the configuration.

The name can contain a maximum of 128 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

Guidelines

The **link-aggregation** command enters Link Aggregation mode to manage the configuration of aggregate interfaces. An aggregate interface can contain only Ethernet interfaces.

Before you can add an Ethernet interface, you must use the **link-aggregation-mode** command in Interface mode to set the Ethernet interface as part of an aggregate interface. While an Ethernet interface is part of the aggregate interface, you cannot change its administrative state.

Use the **no link-aggregation** command to delete the configuration for an aggregate interface.

logging target

This command enters Logging mode.

Syntax

logging target *name*

no logging target *name*

Parameters

name The name of the configuration. The name can have a maximum of 128 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9

- Underscore (_)
- Dash (-)
- Period (.) (note that a name that consists of a single period, or including two periods together, is not permitted)

Guidelines

After you enter Logging mode, you must use the **type** command to identify the log type. More configuration depends on the log type.

Use the **no logging target** command to delete an event log.

packet-capture-advanced

This command manages a packet capture session.

Syntax

Starts a packet-capture session

packet-capture-advanced *name file seconds KB packet-size* ["*expression*"]

Stops a packet-capture session

no packet-capture-advanced *name file*

Parameters

name Specifies the name of the interface for the packet capture. There are two special keywords.

lo Captures internal traffic on the loopback interface.

all Captures traffic from all Ethernet interfaces.

file Specifies the file name for the packet capture. You can simultaneously capture packets on multiple interfaces by specifying a different file name for each capture session or with the special `all` keyword.

seconds
Specifies the maximum duration of the packet-capture session in seconds. Enter a value in the range 5 - 86400. The special value of **-1** indicates that the packet capture is continuous and stops after you enter the **no packet-capture-advanced** command.

KB Specifies the maximum size of the file in KB. Enter a value in the range 10 - 500000.

packet-size
Optionally specifies the maximum size of each captured packet of data in bytes. Enter a value in the range 20 - 9000. The special value **-1** sets the length to the default value of 9000 bytes, which is the maximum transmission unit (MTU) for Ethernet interfaces. The default value ensures that the entire packet is captured.

expression
Optionally specifies the expression that filters the packet capture. Enclose the expression in double quotation marks.

Guidelines

The **packet-capture-advanced** command manages a packet-capture session on the specified interface. The data from the session is saved in the pcap format. To interpret the packet, use a network protocol analyzer.

For a continuous packet capture, you must enter the **no packet-capture-advanced** command to stop the packet capture. When the file for the packet capture reaches its maximum size, a new file is created with the file name as its base file. For example, if the file name is capture-3, the second file is capture-3.002. The appliance retains the last three files.

Use the **no packet-capture-advanced** command to stop a packet-capture session.

Examples

- Start a timed packet-capture session on the eth10 interface that writes data to the temporary:///capture-1 file. Each packet in the capture is limited to the default size of 9000 bytes. The session completes either after 30 minutes or when the file contains 2500 KB, whichever occurs first.

  ```
  # packet-capture-advanced eth10 temporary:///capture-1 1800 2500 -1
  Trace begun.
  #
  ```

- Start a timed packet-capture session on all interfaces that writes data to the temporary:///capture-2 file. The session records only packets where 53 is the destination port. Each packet in the capture is limited to 9000 bytes. The session completes either after 30 minutes or when the file contains 2500 KB, whichever occurs first.

  ```
  # packet-capture-advanced all temporary:///capture-2 1800 2500 9000 "dst port 53"
  Trace begun.
  #
  ```

- Start a continuous packet-capture session on the eth11 interface that writes data to the temporary:///capture-3 file. Each packet in the capture is limited to 500 bytes. The session completes only when you stop it with the **no packet-capture-advanced** command. When the file reaches its maximum size, a new file is created with capture-3 as the base file name; for example, capture-3.002 for the second file that is created.

  ```
  # packet-capture-advanced eth11 temporary:///capture-3 -1 50000 500
  Trace begun.
  #
  ```

- Stop the packet-capture session on the eth11 interface that writes data to the temporary:///capture-3 file.

  ```
  # no packet-capture-advanced eth11 temporary:///capture-3
  Continuous packet capture to temporary:///capture-3 on eth11 stopped.
  #
  ```

raid-activate

This command activates an array volume.

Syntax

raid-activate *name*

Parameters

name Specifies the name of the array volume. The name is raid0.

Guidelines

The **raid-activate** command activates an array volume that is in the `inactive` state, typically with the `foreign volume` inactive state.

Examples

Activate the RAID volume in the disks as the active RAID volume.

```
# raid-activate raid0
```

raid-delete

This command deletes an array volume.

Syntax

raid-delete *name*

Parameters

name Specifies the name of the array volume. The name is `raid0`.

Guidelines

The **raid-delete** command makes the disks that are presently an array volume on the appliance no longer an array volume, removing all metadata. This action destroys the content of the array volume.

Example

Delete the array volume on the disks.

```
# raid-delete raid0
```

raid-learn-battery

This command requests the BBU to start the learning cycle.

Syntax

raid-learn-battery

Guidelines

The **raid-learn-battery** command request the battery backup unit (BBU) to start the learning cycle. This action takes approximately 6 hours to complete. During the learning cycle, the write cache is disabled. During this period, write performance is slower.

Examples

Start the learning cycle for the BBU.

```
# raid-learn-battery
```

raid-make-hot-spare

This command creates a hot spare for a RAID volume.

Syntax

raid-make-hot-spare *name*

Parameters

name Specifies the name of the RAID volume. The name is `raid0`.

Guidelines

The **raid-make-hot-spare** command makes any connected disk that it not part of the RAID volume into a hot spare. A hot spare is the replacement for a failed disk in the RAID volume. Use this command after you replace a failed disk with a new disk.

Example

Create a hot spare for the `raid0` array volume.

```
# raid-make-hot-spare raid0
```

raid-rebuild

This command forces a rebuild of a RAID volume.

Syntax

raid-rebuild *name*

Parameters

name Specifies the name of the RAID volume. The name is `raid0`.

Guidelines

The **raid-rebuild** command forces a rebuild of a RAID volume. The contents of the primary disk are copied to the secondary disk.

Example

Rebuild the `raid0` RAID volume.

```
# raid-rebuild raid0
```

raid-volume-initialize-filesystem

This command initializes the file system of the RAID volume.

Syntax

Initialize the file system without the B2B storage partition.

raid-volume-initialize-filesystem *name*

Initialize the file system with an unencrypted B2B storage partition.

raid-volume-initialize-filesystem *name* **on none**

Initialize the file system with an encrypted B2B storage partition.

raid-volume-initialize-filesystem *name* **on passphrase** *passphrase*

Parameters

name Specifies the name of the RAID volume. The name is `raid0`.

on Initializes the file system and creates the B2B storage partition.

none Indicates that the B2B storage partition is unencrypted.

passphrase
Indicates that the B2B storage partition is encrypted with a passphrase.

passphrase
Specifies the passphrase for encryption.

Guidelines

The **raid-volume-initialize-filesystem** command initializes the file system on the RAID volume to allow it to be made active. When you initialize the file system, the action deletes the existing content.

Examples

- Initializes the file system on the `raid0` RAID volume. The RAID volume has one partition.
  ```
  # raid-volume-initialize-filesystem raid0
  ```
- Initializes the file system on the `raid0` RAID volume. The RAID volume has two partitions. The B2B storage partition is unencrypted.
  ```
  # raid-volume-initialize-filesystem raid0 on none
  ```
- Initializes the file system on the `raid0` RAID volume. The RAID volume has two partitions. The B2B storage partition is encrypted with the `p2ssphras3` passphrase.
  ```
  # raid-volume-initialize-filesystem raid0 on passphrase p2ssphras3
  ```

raid-volume-repair-filesystem

This command repairs the file system for the RAID volume.

Syntax

raid-volume-repair-filesystem *name*

Parameters

name Specifies the name of the RAID volume. The name is `raid0`.

Guidelines

The **raid-volume-repair-filesystem** command repairs the file system on the RAID volume, in case it was corrupted by an abnormal shutdown of the appliance or other error.

Example

Repair the file system on the `raid0` RAID volume.

```
# raid-volume-repair-filesystem raid0
```

raid-volume

This command enters RAID Array mode for a RAID volume.

Syntax

raid-volume *name*

Parameters

name Specifies the name of the array volume. The name is `raid0`.

Guidelines

The **raid-volume** command enters RAID Array mode for the `raid0` RAID volume.

reset username

This command reenables a locked out account.

Syntax

reset username *account* [*password*]

Parameters

account
Specifies the name of the user account to reset.

password
Specifies the new, temporary password for the account.

Guidelines

The **reset username** command allows a privileged administrator to reenable an account after lockout. If the invocation does not include the password, the interface prompts for the password. In either case, the interface prompts for confirmation of the password.

After an administrator reenables the account, the administrator needs to send the owner of the account the new password. The next time the owner of the account logs in, the interface prompts for a new password. This password must comply with the corporate password policy.

Example

Reenable the `suehill` account by changing the password for the account. The administrator does not set the password.

```
# configure terminal
(config)# reset username suehill
Enter new password: ********
Re-enter new password: ********
Password for user 'suehill' is reset.
```

ssh

This command enables SSH on appliance interfaces.

Syntax

ssh *address port*

no ssh [*address*]

Parameters

address
Specifies the IP address of a local interface.

port Identifies the port of a local interface that services SSH traffic. The default value is 22.

Guidelines

SSH is disabled by default. You can use the optional arguments to explicitly bind SSH to a specified interface. If you explicitly bind SSH to an interface, you must have previously configured that interface.

In the absence of an explicit address assignment, SSH first attempts to bind to the management port. If the appliance does not have a management port configured, SSH binds to all configured interfaces.

If the Ethernet for the local address supports IPv6 addresses, modify the `ssh` access control list to include an `allow` clauses for specific or all IPv6 addresses.

Use the **no ssh** command to disable SSH.

Examples

- Enable SSH on port 22 (the default port) of the specified interface.

```
# ssh 10.10.13.4
SSH service listener enabled
```

- Enable SSH on port 2200 of the specified interface.

```
# ssh 10.10.13.4 2200
SSH service listener enabled
```

- Disable SSH on all interfaces, which restores the default state.

```
# no ssh
SSH service listener disabled
```

sslproxy

This command creates an SSL proxy profile that defines an SSL service type.

Syntax

Create an SSL proxy profile for a client

sslproxy *name* **client** *client-profile* [**client-cache** {**on** | **off**}] [**client-sess-timeout** *timer-value*] [**client-cache-size** *entries*] [**permit-insecure-servers** {**on** | **off**}]

sslproxy *name* **forward** *client-profile* [**client-cache** {**on** | **off**}] [**client-sess-timeout** *timer-value*] [**client-cache-size** *entries*] [**permit-insecure-servers** {**on** | **off**}]

Create an SSL proxy profile for a server

sslproxy *name* **server** *server-profile* [**sess-timeout** *timer-value*] [**cache-size** *entries*] [**client-auth-optional** {**on** | **off**}] [**client-auth-always-request** {**on** | **off**}]

sslproxy *name* **reverse** *server-profile* [**sess-timeout** *timer-value*] [**cache-size** *entries*] [**client-auth-optional** {**on** | **off**}] [**client-auth-always-request** {**on** | **off**}]

Create an SSL proxy profile for both client and server

sslproxy *name* **two-way** *server-profile client-profile* [**sess-timeout** *timer-value*] [**cache-size** *entries*] [**client-cache** {**on** | **off**}] [**client-sess-timeout** *timer-value*] [**client-cache-size** *entries*] [**client-auth-optional** {**on** | **off**}] [**client-auth-always-request** {**on** | **off**}] [**permit-insecure-servers** {**on** | **off**}]

sslproxy *name* **both** *server-profile client-profile* [**sess-timeout** *timer-value*] [**cache-size** *entries*] [**client-cache** {**on** | **off**}] [**client-sess-timeout** *timer-value*] [**client-cache-size** *entries*] [**client-auth-optional** {**on** | **off**}] [**client-auth-always-request** {**on** | **off**}] [**permit-insecure-servers** {**on** | **off**}]

Delete an SSL proxy profile

no sslproxy *name*

Parameters

name Specifies the name of the configuration.

The name can contain a maximum of 128 characters. The following characters are valid:

- a through z
- A through Z
- 0 through 9
- Underscore (_)
- Dash (-)
- Period (.)

Note: Names cannot be a single period or two consecutive periods.

client (or forward)
Indicates that the SSL proxy functions as an SSL client in the forward direction. In client proxy mode, SSL is used over the appliance-to-server connection. Client proxy mode requires a client cryptographic profile.

server (or reverse)
Indicates that the SSL proxy functions as an SSL server in the reverse direction. In server proxy mode, SSL is used over the appliance-to-client connection. Server proxy mode requires a server cryptographic profile.

two-way (or both)
Indicates that the SSL proxy functions in both directions as both an SSL client and SSL server. In two-way mode, SSL is used over both the appliance-to-server connection and over the appliance-to-client connection. Two-way mode requires both a client and server cryptographic profile.

server-profile
For the client or two-way operational mode, identifies the cryptographic profile that the SSL client uses to authenticate itself to the SSL server.

client-profile
For the server or two-way operational mode, identifies the cryptographic profile that the SSL server uses to authenticate itself to SSL clients.

sess-timeout *timer-value*
Sets the session timeout value for the server-side session cache in seconds to maintain session-specific state data. Enter a value in the range 1 - 86400. The default value is 300. A value of 0 disables client-side caching.

cache-size *entries*
Optional: Sets the maximum size of the server-side session cache in kilo entries (1024 entries). Enter a value in the range 1 - 500. For example, a value of 10 sets a maximum cache size to 10,240 entries. The default value is 20.

client-cache { **on** | **off** }
Controls client-side caching of session state data. The default value is off.

on Enables client-side caching.

off Disables client-side caching.

client-sess-timeout *timer-value*
Sets the session timeout value for the client-side session cache in seconds to maintain session-specific state data. Enter a value in the range 1 - 86400. The default value is 300. A value of 0 disables client-side caching.

client-cache-size *entries*
Optional: Sets the maximum size of the client-side session cache in entries. Enter a value in the range 1 - 500000. For example, a value of 10000 defines a maximum cache size of 10,000 entries. The default value is 100.

client-auth-optional { **on** | **off** }
When the SSL proxy profile acts as an SSL server, controls when SSL client authentication is optional. The default value is `off`.

on Requests but does not require client authentication. When there is no client certificate, the request does not fail.

off Requires client authentication only when the server cryptographic profile has validation credentials.

client-auth-always-request { **on** | **off** }
When the SSL proxy profile acts as an SSL server, controls when to request SSL client authentication. The default value is `off`.

on Always requests client authentication.

off Requests client authentication only when the server cryptographic profile has validation credentials.

permit-insecure-servers { **on** | **off** }
When the SSL proxy profile acts as an SSL client, specifies whether to permit connections to SSL servers that do not support RFC 5746. The default value is `off`.

on Allows connections to SSL servers that do not support RFC 5746.

off Does not allow connections to SSL server that do not support RFC 5746.

Guidelines

The **sslproxy** command creates an SSL proxy profile that defines an SSL service type: Client, server, or two way. Before you create an SSL proxy profile, you must create one or more cryptographic profiles (Crypto Profile objects). Without cryptographic profiles, the SSL proxy profile is created but in the `down` operational state.

To create a cryptographic profile, use the Crypto **profile** command.

To delete an SSL proxy profile, use the **no sslproxy** command.

Examples

- Create the `SSL-1` server SSL proxy profile with the `Low` cryptographic profile for appliance-to-client connections. Default values are used for the other properties.

  ```
  # sslproxy SSL-1 server Low
  ```

- Create the `SSL-2` client SSL proxy profile with the `High` cryptographic profile for appliance-to-server connections. Default values are used for the other properties.

```
# sslproxy SSL-2 client High
```

- Create the SSL-3 client SSL proxy profile with the ClientIDs cryptographic profile for appliance-to-server connections. Client-side caching is disabled.

  ```
  # sslproxy SSL-3 client ClientIDs client-cache off
  ```

- Create the SSL-4 two-way SSL proxy profile with the NoMD cryptographic profile for appliance-to-client connections and the High cryptographic profile for appliance-to-server connections. Default values are used for the other properties.

  ```
  # sslproxy SSL-4 two-way NoMD High
  ```

- Create the SSL-5 server SSL proxy profile with the Low cryptographic profile for appliance-to-client connection. The session-specific state data times out after 15 minutes (900 seconds), and maximum cache size is allocated for 102,400 entries. Default values are used for the other properties.

  ```
  # sslproxy SSL-5 reverse Low sess-timeout 900 cache-size 100
  ```

- Create the SSL-6 client SSL proxy profile with the High cryptographic profile for appliance-to-server connections. The session-specific state data times out after 15 minutes, and maximum cache size is allocated for 10,000 entries. Default values are used for the other properties.

  ```
  # sslproxy SSL-6 client High client-sess-timeout 900 client-cache-size 10000
  ```

- Create the SSL-7 two-way SSL proxy profile with the NoMD cryptographic profile for appliance-to-client connections and the High cryptographic profile for appliance-to-server connections. The server-side session-specific state data times out after 15 minutes, the maximum server-side cache size is 102,400 entries. The client-side session-specific state data times out after 15 minutes and the maximum client-side cache size is allocated for 10,000 entries. SSL client authentication by the backend server is optional.

  ```
  # sslproxy SSL-7 both NoMD High sess-timeout 900 cache-size 100
     client-sess-timeout 900 client-cache-size 10000
     client-auth-optional on
  ```

- Delete the SSL-7 SSL proxy profile.

  ```
  # no sslproxy SSL-7
  ```

system

This command enters System Settings mode.

Syntax

```
system
```

test hardware

This command tests the hardware.

Syntax

```
test hardware
```

Availability

Physical appliances only.

Guidelines

The **test hardware** command tests the hardware. Depending on the state of the hardware, the command produces output that states the status for each component.

- success
- warning
- failure

The components are broken down into the following categories.
- Backtrace availability
- Interface diagnostics
- Fan diagnostics
- Cryptographic card diagnostics
- RAID volume diagnostics
- Sensors diagnostics
- CPU/memory diagnostics

Samples of success statements are as follows.
- [success] Backtrace file does not exist
- [success] MAC address of interface 'eth10' is 00:11:25:27:bf:e7
- [success] Statistics for interface 'eth10' show no errors
- [success] fan 1 operating within expected range
- [success] Status of the crypto 'standard' is fully operational

Samples of warning statements are as follows.
- [warning] Backtrace file exists.
- [warning] Physical link on interface 'eth10' is down.
- [warning] eth1 has invalid MAC (ff:ff:ff:ff:ff)

Samples of failure statements are as follows.
- [failure] fan 2 operating outside expected range (rpm too low)
- [failure] Status of crypto 'not detected' is unknown.

The output of the test hardware command is included in any generated error report.

top

This command returns you to your initial login mode.

Syntax

top

Guidelines

Regardless of the current location in the configuration modes, the **top** command immediately returns you to your initial login mode.

For custom accounts, this command returns you the your user group-specific login mode.

Examples

Return from Crypto mode to the user-specific login mode.

```
(config crypto)# top
Exiting Crypto Configuration mode
#
```

user-expire-password

This command forces a user to change the account password at the next log on.

Syntax

user-expire-password *account*

Parameters

account
Identifies the target account.

Example

Force a password change for the `josephb` account on the next log on.

```
# user-expire-password josephb
Expire password for user 'josephb' succeeded
```

write memory

This command copies the running configuration as the startup configuration.

Syntax

write memory

Guidelines

After the running configuration is copied to the `config:///autoconfig.cfg` file, the appliance determines whether the current startup configuration script file can be overridden by `config:///autoconfig.cfg`.

If it can be overridden, the `autoconfig.cfg` file becomes the startup configuration.

Examples

- Save the running configuration as the startup configuration.

  ```
  # write memory
  Overwrite existing autoconfig.cfg? y
  #
  ```

- Cancel the operation.

  ```
  # write memory
  Overwrite existing autoconfig.cfg? n
  #
  ```

Host Alias commands

Host Alias mode provides the commands to create or modify a host alias.

To enter the configuration mode, use the Global **host-alias** command. To delete a host alias, use the Global **no host-alias** command.

While in this mode, use the commands in the following table to define a host alias.

- To view the current configuration, use the **show** command.
- To restore default values, use the **reset** command.

- To exit this configuration mode without saving changes to the running configuration, use the **cancel** command.
- To exit this configuration mode and save changes to the running configuration, use the **exit** command.

Table 35. Host Alias commands

Command	Purpose
"**admin-state**" on page 437	Sets the administrative state for the configuration.
"**ip-address**"	Creates an alias for an IP address on a Ethernet port.

ip-address

This command creates an alias for an IP address on a Ethernet port.

Syntax

ip-address *address*

Parameters

address
Specifies the IP address to map.

Guidelines

The **ip-address** command creates an alias for a local IP address of the appliance. Instead of providing the IP address, you can specify this alias.

Example

Create the `Ragnarok` alias. Map `Ragnarok` to IP address 192.168.12.12.

```
# host-alias Ragnarok
New Host Alias configuration
# ip-address 192.168.12.12
# exit
#
```

Initialization commands

You can use the initialization commands to initalize or reset the IBM MQ Appliance.

The initialization commands can be run from the command line interface in flash mode. To enter flash configuration mode, from the appliance command line, type `flash`.

boot config

This command designates the startup configuration for the next restart.

Syntax

boot config *file*

Parameters

file Specifies the name of the startup configuration file.

Guidelines

The **boot config** and **boot image** commands work together to define the restart process.

- The **boot config** command designates the startup configuration.
- The **boot image** command designates the startup firmware image.

Example

Designate `testEnvironment.cfg` as the startup configuration.

```
# boot config testEnvironment.cfg
#
```

boot delete

This command deletes the secondary installation.

Syntax

boot delete

Guidelines

A firmware upgrade with the **boot image** command retains current configuration data to restore the appliance to a known, stable state if necessary. The previous firmware image and associated configuration data is referred to as the *secondary* installation.

While you can use the **boot delete** command to delete the secondary installation, keep in mind that its deletion prevents firmware rollback as provided by the **boot switch** command. Therefore, do not delete the secondary installation unless directed by IBM Support.

Example

Delete the secondary installation.

```
# boot delete
Previous firmware install deleted
```

boot image

This command designates the startup firmware image and restarts the appliance with this image.

Syntax

boot image [**accept-license**] *file*

Parameters

accept-license
: Indicates acceptance of the terms of the license agreements.

file
: Specifies the name of the firmware image.

Guidelines

The **boot config** and **boot image** commands work together to define the restart process.

- The **boot config** command designates the startup configuration.
- The **boot image** command designates the startup firmware image.

The firmware image can contain new version of component firmware. Examples of component firmware includes, but is not limited to, the BIOS, BMC, and RAID controller for the appliance. When the firmware image contains new versions of component firmware. the upgrade can take approximately 20 minutes. During this upgrade process, do not power off or restart the appliance.

You must use the `accept-license` keyword to confirm acceptance of the original license agreement (that is, the license agreement that you accepted when you installed and configured the IBM MQ Appliance).

You can use an empty `license.accepted` file when you automate the installation process with a script. To indicate acceptance of the terms of the license agreements, the script must create the empty `license.accepted` file in the `temporary:` directory on the appliance. When the file exists on the appliance:

- The license is considered to be accepted.
- The appliance allows the firmware upgrade to proceed.

Examples

- Restarts the appliance with the file `5001` in the `image:` directory. Acknowledge acceptance of the license agreements with the **accept-license** keyword.

```
# boot image accept-license 5001
...
.....Firmware upgrade successful
Device is rebooting now.
```

- Restarts the appliance with the file `5001` in the `image:` directory when an empty file `temporary://license.accepted` exists.

```
# boot image 5001
...
.....Firmware upgrade successful
Device is rebooting now.
```

boot switch

This command switches between primary and secondary installations.

Syntax

boot switch

Guidelines

A firmware upgrade with the **boot image** command retains current configuration data, including activated features, which allows you to restore the appliance (rolled back) to a known, stable state if necessary.

- The previous firmware image and associated configuration data is the *secondary* installation.
- The newly installed firmware image and associated configuration data is the *primary* installation.

The **boot switch** command is not available on an appliance that is part of a high availability group.

When you switch between firmware images that contain different versions of component firmware, the switch can take approximately 20 minutes. During this switch operation, do not power off or restart the appliance. Examples of component firmware includes, but is not limited to, the BIO, BMC or RAID controller for the appliance.

You can use the **boot switch** command to transition between the newly installed and previously active firmware versions.

Configuration edits that are made to a selected image are local and are not included during the rollback operation. It is possible to issue the **boot switch** command twice, which returns the appliance to the newly installed version.

While you can use the **boot delete** command to delete the secondary installation, keep in mind that its deletion prevents firmware rollback as provided by the **boot switch** command. Do not delete of the secondary installation unless you are working with IBM Support.

Example

The rollback operation failed. The secondary installation was deleted with the **boot delete** command.

```
# boot switch
% Firmware roll-back failed: Switch active firmware failed:
Secondary install not available
```

boot update

This command creates a new configuration or opens an existing configuration for editing.

Syntax

boot update write *file*

boot update append *file*

Parameters

write Creates and opens a new configuration. If the file exists, the appliance erases and opens the existing configuration.

append Opens an existing configuration to which to add commands.

file Specifies the name of the configuration file to create or open.

Guidelines

The **boot update** command creates a new configuration or opens an existing configuration for editing. When you open a configuration file, the CLI prompts for input. `Enter startup commands, one per line. End with a period.`

Enter commands, and press the Enter key.

If you add commands to an existing configuration, start with appropriate commands to move to the correct configuration mode.

Follow the last command of the configuration with the following sequence to signal the end of the configuration.

1. Press the Enter key.
2. Enter a `..`
3. Press the Enter key.

The CLI acknowledges configuration completion. `Configuration completed successfully.`

After you complete your edits, use the **`boot config`** and **`shutdown`** commands to activate this configuration.

Examples

- Create the `jrb_03.cfg` configuration. If it exists, the appliance erases and opens the file.

```
# boot update write jrb_03.cfg
Enter startup commands, one per line. End with a period.
```

- Open the `jrb_03.cfg` configuration to add commands.

```
# boot update append jrb_01.cfg
Enter startup commands, one per line. End with a period.
```

`reinitialize`

This command deletes all configuration data from the file system of the appliance.

Syntax

`reinitialize` *file*

Parameters

file Specifies the name of the firmware image to reinitialize the appliance. The file must be in the `image:` directory.

Guidelines

The **`reinitialize`** command deletes all configuration data in the file system of the appliance. This data consists of style sheets, object configurations, keys, certificates, log files, and so forth.

After files are deleted, they cannot be recovered. If you might need any of these files after you reinitialize the appliance, ensure that you have copies of these files.

The following conditions occur after you reinitialize the appliance.

- The network configuration is removed. You can no longer access the appliance through the former IP address. You can connect to the appliance only through a serial cable. After you connect to the appliance, use the `startup` command to provide the base configuration data.
- The password for the `admin` account is reset to `admin`.

Example

Delete all user files and data on the appliance, and restart the appliance.

```
# reinitialize firmware.scrypt2
WARNING - all user data and files will be deleted
Do you want to continue ("yes" or "no")? y
#
```

shutdown

This command restarts or shuts down the appliance.

Syntax

shutdown

shutdown reboot [*seconds*]

shutdown reload [*seconds*]

shutdown halt [*seconds*]

Parameters

reboot Shuts down and restarts the appliance.

reload Restarts the appliance.

halt Shuts down the appliance and, by default, restarts after 20 minutes.

seconds
Specifies the number of seconds before the appliance starts the shutdown operation. Enter a value in the range 0 - 65535. The default value is 10.

Guidelines

The **shutdown** command shuts down, restarts, or shuts down and restarts the appliance. Without parameters, the command restarts the appliance.

The appliance restarts with the startup configuration that is specified by the **boot config** command and the startup firmware image that is specified by the **boot image** command. Without a designated startup configuration or firmware image, the appliance restarts with the configuration and firmware image that were active when you issued the **shutdown** command.

Examples

- Wait 10 seconds to shut down and restart the appliance.

```
# shutdown reboot
Reboot in 10 second(s).
#
```

- Wait 20 seconds to restart the appliance.

```
# shutdown reload 20
Reload in 20 second(s).
#
```

- Wait 1 minute to shut down the appliance. If you do not restart or power down the appliance, it restarts in 20 minutes.

```
# shutdown halt 60
Shutdown in 60 second(s).
#
```

IPMI user commands

You can use the IPMI user commands to create or modify an IPMI users on the IBM MQ Appliance.

The IPMI user commands can be run from the command line interface in IPMI user configuration mode. To enter IPMI user configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:
   `config`
2. From global configuration mode, enter IPMI user configuration mode:
   `ipmi-user`
3. Type `exit` to save your changes and leave IPMI user configuration mode, then type `exit` again to leave global configuration mode.

password

This command sets the password for the IPMI user.

Syntax

password [*password*]

Parameters

password
Optional: Sets the password for the user.

Guidelines

The **password** command sets the password that the remote user must present for authentication. The password length must be 8 - 16 characters. This command with the **user-id** command defines an authentication record in the Baseboard Management Controller BMC).

Note: Because the password is in only the BMC, it is not included as part of an export or backup operation. Because the password is not part of an export, it is not added during an import or restore operation.

user-id

This command sets the identifier for the IPMI user.

Syntax

user-id *identifier*

Parameters

identifier
Sets the identifier for the IPMI user. Enter a value in the range 3 - 10. The default value is 3.

Guidelines

The **user-id** command sets the identifier for the IPMI user. This command with the **password** command defines an authentication record in the Baseboard Management Controller BMC).

Each user must have a unique identifier. The index for all user configurations in the BMC is this identifier.

IPMI LAN channel commands

You can use the IPMI LAN channel commands to create or modify an IPMI LAN channel on the IBM MQ Appliance.

The IPMI LAN channel commands can be run from the command line interface in IPMI LAN channel configuration mode. To enter IPMI LAN channel configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:
   `config`
2. From global configuration mode, enter IPMI LAN channel configuration mode:
   `ipmi-lan-channel`
3. Type `exit` to save your changes and leave IPMI LAN channel configuration mode, then type `exit` again to leave global configuration mode.

allowed-user

This command specifies the users allowed to use the channel.

Syntax

Allows a user

allowed-user *ID* [*privilege*] [*SOL*] [*sessions*]

Disallows a user

no allowed-user *ID*

Parameters

ID Specifies the name of an existing Intelligent Platform Management Interface (IPMI) user to enable on this LAN channel. To create an IPMI user, use the Global **ipmi-user** command.

privilege

Optional: Sets the maximum privilege level to allow the user on this LAN channel.

callback
Allows IPMI commands that establish a callback connection.

user Allows IPMI commands that read operational status. This setting is the default value.

operator
Allows IPMI commands that change operational status.

If this value is greater than the value set with the **maximum-channel-privilege-level** command, the value set with the **maximum-channel-privilege-level** command is the effective level.

SOL Optional: Indicates whether to allow the user to use the serial over LAN protocol (payload) over Remote Control and Management Protocol Plus (RCMP+) on this LAN channel.

on The user can use serial over LAN provided that the user has a privilege level equivalent to or greater than the minimum required privilege level set with the **sol-required-user-privilege-level** command. This setting is the default value.

off The user cannot use serial over LAN.

sessions

Optional: Sets the maximum number of simultaneous sessions to allow the user on this LAN channel. The minimum value is 0, and the maximum value is 7. If set to 0, there is no limit. The default value is 0.

Guidelines

The **allowed-user** command specifies a user allowed to use this LAN channel. For each user, include the maximum privilege level, whether to allow serial over LAN, and the maximum number of simultaneous sessions. Use this command for each user to allow.

CAUTION:
If you enable serial over LAN support for the IPMI LAN channel, an IPMI user who connects through an IPMI 2.0 client has access to the serial console port.

- **An IPMI user can connect through an IPMI 2.0 client independent of the state of the appliance. The only time that an IPMI user cannot connect to the appliance is when the appliance is disconnected from AC power.**
- **An IPMI user has higher priority than a user who is directly connected to the serial port.**
- **If a user is directly connected to the serial port, the IPMI user usurps the current serial session and does not need to log in to the appliance. If no user is directly connected to the serial port, the IPMI user must log in to the appliance.**
- **There can be only one serial user. The IPMI user suspends the serial session of the user who is directly connected to the serial port until the remote IPMI user closes (deactivates) the serial session. When the IPMI user closes the serial session, the session for the user who is directly connected to the serial port resumes.**
- **If an IPMI user is connected to the appliance and you need to use the serial port, you must unplug the Ethernet cable from the MGT0 port. After you unplug the cable, wait for the serial port to become available, which can take up to 20 minutes.**

Use the **no allowed-user** command to disallow a user on this LAN channel.

ip address

This command sets IP addresses with subnet mask for the IPMI LAN channel.

Syntax

Assign the address.

ip address *address*

Remove the address.

no ip address *address*

Parameters

address

Specifies the IP version 4 (IPv4) address and netmask. The netmask can be in CIDR (slash) format or dotted decimal format. With CIDR format, the integer specifies the prefix length. The prefix length can be in the range 0 - 32.

Guidelines

The **ip address** command sets IPv4 addresses with subnet mask to the interface. This IP address must be distinct from all IP addresses on the appliance and from all IP address on the connected subnet (broadcast domain).

Use the **no ip address** command to remove the address.

ip default-gateway

This command sets the default gateway for the IPMI LAN channel.

Syntax

Sets the default gateway

ip default-gateway *address*

Deletes the default gateway

no ip default-gateway

Parameters

address

Specifies the IP address of the default gateway.

Guidelines

The **ip default-gateway** command designates the default gateway that communicates with systems not on the local subnet. The default gateway must be on the same subnet as the IP address defined with the **ip address** command. Without a default gateway, the LAN channel cannot communicate with clients that are not on the same local subnet (broadcast domain).

Use the **no ip default-gateway** command to delete the default gateway.

maximum-channel-privilege-level

This command sets the maximum privilege level for users.

Syntax

maximum-channel-privilege-level *privilege*

Parameters

privilege

Sets the maximum privilege level for users on this channel. The keyword list is in from lowest to highest privilege level.

callback

Allows IPMI commands that establish a callback connection.

user Allows IPMI commands that read operational status. This setting is the default value.

operator

Allows IPMI commands that change operational status.

Guidelines

The **maximum-channel-privilege-level** command sets the maximum privilege level for all users on the LAN channel.

Each Intelligent Platform Management Interface (IPMI) command has a minimum privilege level. This command sets the maximum privilege level that a user can attain on the LAN channel even when that user has a greater privilege level.

sol-enabled

This command indicates whether to support serial over LAN.

Syntax

sol-enabled { **on** | **off** }

Parameters

on Users can use the serial over LAN protocol. This setting is the default value.

off Users cannot use the serial over LAN protocol.

Guidelines

The **sol-enabled** command indicates whether to support the serial over LAN protocol over Remote Control and Management Protocol Plus (RCMP+) on this channel. When enabled, use the **sol-required-user-privilege-level** command to set the privilege level for serial over LAN connections.

CAUTION:
If you enable serial over LAN support for the IPMI LAN channel, an IPMI user who connects through an IPMI 2.0 client has access to the serial console port.

- **An IPMI user can connect through an IPMI 2.0 client independent of the state of the appliance. The only time that an IPMI user cannot connect to the appliance is when the appliance is disconnected from AC power.**
- **An IPMI user has higher priority than a user who is directly connected to the serial port.**
- **If a user is directly connected to the serial port, the IPMI user usurps the current serial session and does not need to log in to the appliance. If no user is directly connected to the serial port, the IPMI user must log in to the appliance.**
- **There can be only one serial user. The IPMI user suspends the serial session of the user who is directly connected to the serial port until the remote IPMI user closes (deactivates) the serial session. When the IPMI user closes the serial session, the session for the user who is directly connected to the serial port resumes.**
- **If an IPMI user is connected to the appliance and you need to use the serial port, you must unplug the Ethernet cable from the MGT0 port. After you unplug the cable, wait for the serial port to become available, which can take up to 20 minutes.**

sol-required-user-privilege-level

This command sets the privilege level for serial over LAN.

Syntax

sol-required-user-privilege-level *privilege*

Parameters

privilege
Sets the privilege level required for serial over LAN. The keyword list is in from lowest to highest privilege level.

callback
Allows IPMI commands that establish a callback connection.

user Allows IPMI commands that read operational status. This setting is the default value.

operator
Allows IPMI commands that change operational status.

Guidelines

The **sol-required-user-privilege-level** command sets the required privilege level for users to use the serial over LAN protocol (payload) over Remote Control and Management Protocol Plus (RCMP+) on this LAN channel.

This command is available when serial over LAN is enabled with the **sol-enabled** command.

Link aggregation commands

You can use the link aggregation commands to configure link aggregation interfaces on the IBM MQ Appliance.

The link aggregation commands can be run from the command line interface in link aggregation configuration mode. To enter link aggregation configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:
   `config`
2. From global configuration mode, enter link aggregation configuration mode:
   `link-aggregation` *name*

   where *name* is the name of the link aggregation interface that you want to configure. The name can have a maximum of 128 characters. The following characters are valid:
   - a through z
   - A through Z
   - 0 through 9
   - Underscore (_)
   - Dash (-)
   - Period (.) (note that a name comprising a single period, or including two periods together, is not permitted)
3. Type `exit` to leave the configuration mode and save your changes, then type `exit` again to leave global configuration mode.

Link aggregation interfaces are not supported for links used in high availability configurations or disaster recovery configurations.

ip-address

This command assigns the primary network address for the aggregate interface.

Syntax

ip-address *address*

Parameters

address
Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

Guidelines

The **ip-address** command assigns the primary network address to the interface. The network address is an IP address with its subnet mask.

To assign secondary, or auxiliary, IP addresses, use the **ip-secondary-address** command.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

Examples

- Assign an IP address in version 4 format.

```
# ip-address 192.168.7.6/27
```

- Assign an IP address in version 6 format.

```
# ip-address 2001:0db8:3c4d:0015::abcd:ef12/34
```

ip-config-mode

This command identifies the configuration mode for the aggregate interface.

Syntax

ip-config-mode { **static** | **dhcp** | **slaac** }

Parameters

static Indicates a static, manual configuration. This setting is the default value.

dhcp Indicates IPv4 autoconfiguration with DHCP.

slaac Indicates IPv6 autoconfiguration with SLAAC.

Guidelines

The **ip-config-mode** command identifies the configuration mode of the interface.

- With the `static` keyword, define the configuration for the interface as provided by your network administrator.
  - Use the **ip-address** command to assign the primary network address.
  - Use the **ip-secondary-address** command to manage secondary, or auxiliary, network addresses.
  - Use the **ipv4-default-gateway** command to assign the default IPv4 gateway.
  - Use the **ipv6-default-gateway** command to assign the default IPv6 gateway.
  - Use the **ip-route** command to manage static routes in the routing table.

- With the `dhcp` keyword, the appliance ignores configuration data about the physical interface.
- With the `slaac` keyword, the appliance ignores configuration data about the physical interface.

Examples

- Change the configuration mode to IPv4 autoconfiguration with DHCP.

  ```
  # ip-config-mode dhcp
  ```

- Change the configuration mode to manual configuration.

  ```
  # ip-config-mode static
  ```

ip-route

This command manages static routes in the routing table for the aggregate interface.

Syntax

Add a static route

ip-route *address next-hop-address* [*metric*]

Delete a static route

no ip-route *address next-hop-address*

Parameters

address

Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

next-hop-address

Specifies the IP address of the next-hop router.

metric Optionally specifies the preference for the route. The lesser the value, the more preferred the route. For each IP family, the supported range differs.

- For IPv4, enter a value in the range 0 - 255. The default value is 0.
- For IPv6, enter a value in the range 0 - 65536. The default value is 512.

Guidelines

The **ip-route** command manages static routes in the routing table. Issue this command for each static route to add to the routing table.

To delete a static route, use the **no ip-route** command. Issue this command for each static route to delete from the routing table.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

Examples

- Add a static route to the routing table (subnet 10.10.10.224 via next-hop router 192.168.1.100). The metric for the route is 0, the default value for IPv4, which is the most preferred route.

  ```
  # ip-route 10.10.10.0/27 192.168.1.100
  ```

- Delete a static route from the routing table (subnet 10.10.10.224 via next-hop router 192.168.1.100).

  ```
  # no ip-route 10.10.10.0/27 192.168.1.100
  ```

ip-secondary-address

This command manages secondary network addresses for the aggregate interface.

Syntax

Add a secondary address

ip-secondary-address *address*

Remove a secondary address

no ip-secondary-address *address*

Remove all secondary addresses

no ip-secondary-address

Parameters

address

Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

Guidelines

The **ip-secondary-address** command manages secondary network addresses for the current interface. The network address is the IP address and its subnet mask. A secondary IP address is a bind address. The secondary IP address is used only as a source IP address when it responds to incoming traffic to the secondary IP address.

To create the primary IP address, use the **ip-address** command.

To remove secondary IP addresses, use the **no ip-secondary-address** command.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

Examples

- Add 192.168.7.6/27 as a secondary IP address to the interface.

  ```
  # ip-secondary-address 192.168.7.6/27
  ```

- Remove 192.168.7.6/27 as a secondary IP address.

  ```
  # no ip-secondary-address 192.168.7.6/27
  ```

- Remove all secondary IP addresses.

  ```
  # no ip-secondary-address
  ```

ipv4-default-gateway

This command designates the default IPv4 gateway for the aggregate interface.

Syntax

Designates the default IPv4 gateway

ipv4-default-gateway *address*

Deletes the default IPv4 gateway
no ipv4-default-gateway

Parameters

address
Specifies the IP address of the default IPv4 gateway.

Guidelines

The **ipv4-default-gateway** command designates the default IPv4 gateway that the interface can reach. If the interface supports both IP families, use the **ipv6-default-gateway** command to designate the default IPv6 gateway.

Use the **no ipv4-default-gateway** command to delete the default IPv4 gateway.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

ipv6-dadtransmits

This command sets the number of IPv6 duplication address detection attempts for the aggregate interface.

Syntax

ipv6-dadtransmits *attempts*

Parameters

attempts
Specifies the number of attempts. The default value is 1.

Guidelines

The **ipv6-dadtransmits** command sets the number of IPv6 duplication address detection (DAD) attempts. This command is relevant for only IPv6 addresses on the appliance.

If you specify more than one attempt, use the **ipv6-nd-retransmit-timer** command to set the interval between attempts.

ipv6-default-gateway

This command designates the default IPv6 gateway for the aggregate interface.

Syntax

Designate the default IPv6 gateway
ipv6-default-gateway *address*

Delete the default IPv6 gateway
no ipv6-default-gateway

Parameters

address
Specifies the IP address of the default IPv6 gateway.

Guidelines

The **ipv6-default-gateway** command designates the default IPv6 gateway that the interface can reach. Define a default IPv6 gateway if you defined IPv6 IP addresses.

If the interface supports both IP families, use the **ipv4-default-gateway** command to designate the default IPv4 gateway.

Use the **no ipv6-default-gateway** command to delete the default IPv6 gateway.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

ipv6-nd-retransmit-timer

This command sets the interval between IPv6 neighbor discovery attempts for the aggregate interface.

Syntax

ipv6-nd-retransmit-timer *milliseconds*

Parameters

milliseconds
Specifies the interval between attempts in milliseconds. The default value is 1000.

Guidelines

The **ipv6-nd-retransmit-timer** command sets the interval neighbor discovery attempts. This command is relevant for only when the interface uses IPv6 addresses.

lacp-hash

This command sets which hash function the LACP aggregation uses to determine the interface for outbound packets.

Syntax

lacp-hash { **ip-port** | **mac** | **mac-ip** }

Parameters

ip-port
Indicates that the hash function uses both IP addresses and TCP/UDP ports. IP addresses are layer 3. TCP/UDP ports are layer 4. This approach is not strictly compliant to the LACP standard.

mac Indicates that the hash function uses only MAC addresses, which is layer 2 data.

mac-ip Indicates that the hash function uses both MAC addresses and IP addresses. MAC addresses are layer 2. IP addresses are layer 3. This setting is the default value.

Guidelines

The **lacp-hash** command sets which hash function the Link Aggregation Control Protocol (LACP) aggregation uses to determine the interface for outbound packets.

The command is meaningful only when you use the **type** command to define LACP as the link aggregation type.

Example

Set the aggregate interface to use LACP and choose the aggregator with the highest bandwidth. The aggregate interface uses only the MAC address hash algorithm.

```
# type LACP
# lacp-select bandwith
# lacp-hash mac
```

lacp-select

This command sets the algorithm for the LACP selection policy.

Syntax

lacp-select { **bandwith** | **count** | **stable** }

Parameters

bandwith
Chooses the aggregator with the highest bandwidth.

count Selects the aggregator with the most NICs.

stable Does not change the aggregator when a better one is available. This setting is the default value.

Guidelines

The **lacp-select** command sets the algorithm for the Link Aggregation Control Protocol (LACP) selection policy. When there is more than one LACP aggregator for the members of an LACP aggregation, the appliance uses the algorithm to determine which aggregator to use. In other words, the algorithm determines which group of aggregate interfaces is active.

The command is meaningful only when you use the **type** command to define LACP as the link aggregation type.

Example

Set the aggregate interface to use LACP and choose the aggregator with the highest bandwidth. The aggregate interface uses only the MAC address hash algorithm.

```
# type LACP
# lacp-select bandwith
# lacp-hash mac
```

link

This command manages Ethernet interfaces in the aggregate interface.

Syntax

Add an Ethernet interface
link *name*

Delete an Ethernet interface
no link *name*

Parameters

name Specifies the name of an Ethernet interface.

Guidelines

The **link** command manages Ethernet interfaces in the aggregate interface.

- For each Ethernet interface to add to the aggregate interface, use the **link** command.
- For each Ethernet interface to delete from the aggregate interface, use the **no link** command.

You cannot add members to an aggregate interface in the following situations:

- The Ethernet interface is part of a VLAN interface
- The Ethernet interface is part of another aggregate interface
- The Ethernet interface is not enabled for link aggregation with the **link-aggregate-mode** command in Ethernet mode.

Examples

Add the `eth10`, `eth11`, and `eth12` Ethernet interfaces to the aggregate interface.

```
# link eth10
# link eth11
# link eth12
```

Deletes the `eth12` Ethernet interface from the aggregate interface.

```
# no link eth12
```

packet-capture

This command manages a packet-capture for the aggregate interface session.

Syntax

Start a packet-capture session
packet-capture *file seconds KB* ["*expression*"]

Stop a packet-capture session
no packet-capture *file*

Parameters

file Specifies the file name for the packet capture. You can simultaneously capture packets on multiple interfaces by specifying a different file name for each interface.

seconds
Specifies the maximum duration of the packet-capture session in seconds. Enter a value in the range 5 - 86400. The special value of `-1` indicates that the packet capture is continuous and completes when it reaches the maximum file size or until you issue the **no packet-capture** command.

KB Specifies the maximum size of the file in KB. Enter a value in the range 10 - 500000.

expression
Optionally specifies the expression that filters the packet capture. Enclose the expression in double quotation marks.

Guidelines

The **packet-capture** command manages a packet-capture session on the current interface. The data from the session is saved in the `pcap` format. To interpret the packet, use a network protocol analyzer.

Use the **no packet-capture** command to stop a packet-capture session.

Examples

- Start a timed packet-capture session that writes data to the `temporary:///capture-1` file. The session completes either after 30 minutes or when the file contains 2500 KB, whichever occurs first.

```
# packet-capture temporary:///capture-1 1800 2500
Trace begun.
#
```

- Start a timed packet-capture session that writes data to the `temporary:///capture-2` file. The session records only packets where 53 is the destination port. The session completes either after 30 minutes or when the file contains 2500 KB, whichever occurs first.

```
# packet-capture temporary:///capture-2 1800 2500 "dst port 53"
Trace begun.
#
```

- Start a continuous packet-capture session that writes data to the `temporary:///capture-3` file. The session completes either when it contains 50000 KB or when you stop it.

```
# packet-capture temporary:///capture-3 -1 50000
Trace begun.
#
```

- Stop the packet-capture session that writes data to the `temporary:///capture-3` file.

```
# packet-capture temporary:///capture-3
Continuous packet capture to temporary:///capture-3 on interface stopped.
#
```

type

This command defines the mode to use for link aggregation.

Syntax

type { **active-backup** | **LACP** | **transmit** }

Parameters

active-backup
Sets link aggregation to active-backup. One link is active, and the other links are backup. If the active link is lost, switches to a backup link. This setting is the default value.

LACP Sets link aggregation to the Link Aggregation Control Protocol (LACP). The IEEE 802.1AX-2008 standard defines LACP. This mode requires support on the network switch.

transmit
Sets link aggregation to transmit-based load balancing. This mode uses a single link for incoming traffic but distributes outgoing traffic among all links.

Guidelines

The **type** command defines the mode to use for link aggregation.

When the aggregate interface uses LACP, you can change the selection policy and the distribution algorithm for outbound packets.

- Use the **lacp-select** command to change the LACP selection policy. The default selection policy is to use active-backup.
- Use the **lacp-hash** command to change the LACP distribution algorithm for outbound packets. The default algorithm is to use the hash function against both MAC addresses and IP addresses.

Example

Set the aggregate interface to use LACP and choose the aggregator with the highest bandwidth. The aggregate interface uses only the MAC address hash algorithm.

```
# type LACP
# lacp-select bandwith
# lacp-hash mac
```

Log target commands

Log target mode provides the commands to create or modify a log target.

To enter the mode, use the Global **logging target** command. To delete a log target, use the **no logging target** command.

While in this mode, use the commands in the following table to configure a log target.

- To view the current configuration, use the **show** command.
- To restore default values, use the **reset** command.
- To exit this configuration mode without saving changes to the running configuration, use the **cancel** command.
- To exit this configuration mode and save changes to the running configuration, use the **exit** command.

Table 36. Log Target commands

Command	Purpose
"**active-timeout**" on page 553	This command sets the timer to close an established and active connection to the server.
"**admin-state**" on page 437	This command sets the administrative state for the configuration.
"**ansi-color**" on page 553	This command manages a multicolored console log.
"**archive-mode**" on page 553	This command specifies an archival behavior for file-based logs.
"**backup**" on page 554	This command specifies a backup for the current log.

Table 36. Log Target commands (continued)

Command	Purpose
"**connect-timeout**" on page 554	This command sets the time to wait to establish a connection.
"**email-address**" on page 555	This command specifies the email address of a remote recipient of SMTP log messages.
"**event**" on page 555	This command adds an event class and a priority to the log target.
"**event-code**" on page 556	This command specifies an event code to add to the log target.
"**event-detection**" on page 556	This command suppresses identical events in the log target.
"**event-filter**" on page 557	This command specifies an event code to exclude from the log target.
"**facility**" on page 557	This command specifies the syslog facility.
"**feedback-detection**" on page 558	This command suppresses events from the logging subsystem itself.
"**format**" on page 558	This command specifies the format in which events are added to the log.
"**idle-timeout**" on page 559	This command sets the idle timer for the logging target.
"**ip-filter**" on page 559	This command adds an IP address to include in the log target.
"**local-address**" on page 560	This command specifies the local address over which log events are transmitted to a remote recipient.
"**local-file**" on page 560	This command specifies a local file for log messages.
"**local-ident**" on page 560	This command sets the string that the remote recipient uses to identify this log target.
"**object**" on page 561	This command adds an object filter to the current log target.
"**priority**" on page 561	This command sets the service-level priority for the log target.
"**rate-limit**" on page 562	This command sets the maximum rate of events to log.
"**remote-address**" on page 562	This command specifies the destination address of log messages or the log itself.
"**remote-directory**" on page 563	This command specifies the remote directory where uploaded logs are stored.
"**remote-login**" on page 564	This command specifies the user name to use to upload a log file to a remote server.
"**remote-port**" on page 565	This command specifies the listening port on the remote server.
"**rotate**" on page 565	This command sets the maximum number of file rotations.
"**sender-address**" on page 566	This command specifies the email address of the sender.
"**size**" on page 566	This command sets the maximum size of a local log file.
"**smtp-domain**" on page 567	This command specifies the domain name of the SMTP client.
"**soap-version**" on page 567	This command specifies the version of SOAP to use.
"**ssl**" on page 568	This command associates an SSL proxy profile for SOAP-based log over HTTPS.
"**suppression-period**" on page 568	This command sets the interval to suppress identical events.
"**timestamp**" on page 568	This command specifies the time stamp format.
"**trigger**" on page 568	This command creates an event trigger.
"**type**" on page 569	This command identifies the logging model.
"**upload-method**" on page 571	This command sets the protocol to upload a file-based log to a remote site.
"**url**" on page 572	This command sets the destination for SOAP-based log entries.

active-timeout

This command sets the timer to close an established and active connection to the server.

Syntax

active-timeout *seconds*

Parameters

seconds

Sets the number of seconds to wait before the connection is closed. Enter a value in the range 0 - 60. The default value is 0.

Guidelines

The **active-timeout** command sets the number of seconds to wait before an established and active connection to the server is closed. The default value of 0 allows the log target to most efficiently send messages to the server by maintaining a working connection indefinitely.

Attention: If multiple log targets have the following configuration, they might share connections.

- The same local address and port
- The same remote address and port

Because of potential connection-sharing, set the same active timeout for these log targets.

You can use the **connect-timeout** and **idle-timeout** connection commands to modify the additional timeout values.

- The connect timeout sets the number of seconds to wait to establish a connection to the server before an error message is generated.
- The idle timeout sets the number of seconds to wait before an established, but inactive, connection to the server is closed.

ansi-color

This command manages a multicolored console log.

Syntax

ansi-color { **on** | **off** }

Parameters

on Enables different priorities to display in different colors.

off Provides a monochrome display. This setting is the default value.

Guidelines

Meaningful only when the log type is `console`. Otherwise, it is ignored.

archive-mode

This command specifies an archival behavior for file-based logs.

Syntax

archive-mode { **rotate** | **upload** }

Parameters

rotate Specifies that when a log file reaches its maximum size, the log is rotated as specified by the **rotate** command. This setting is the default value.

upload Specifies that when a log file reaches its maximum size, the file is uploaded to a specified site for remote storage.

Guidelines

The **archive-mode** command is required when the log type is `file`. Otherwise, it is ignored.

After you set the upload mode, you must use the **remote-address**, **remote-directory**, **remote-login**, and **upload-method** commands to enable transfer of the log file to the remote site.

Examples

- Set the archive type to upload.

  ```
  # archive-mode upload
  ```

- Set the archive type to rotate, which restores the default state.

  ```
  # archive-mode rotate
  ```

connect-timeout

This command sets the time to wait to establish a connection.

Syntax

connect-timeout *seconds*

Parameters

seconds

Sets the number of seconds that the appliance waits to establish a connection. Enter a value in the range 1 - 90. The default value is 60.

Guidelines

The **connect-timeout** command indicates the time that the appliance waits to establish a connection to the server before it generates an error message. After the appliance generates the log message, it attempts to establish the connection again. You can use the **active-timeout** and **idle-timeout** connection commands to modify other timeout values.

- The active timer sets the number of seconds to wait before the appliance closes an established and active connection to the server.
- The idle timer sets the number of seconds to wait before the appliance closes an established but inactive connection to the server.

backup

This command specifies a backup for the current log.

Syntax

backup *name*

Parameters

name Specifies the name of a log, of any log type.

Guidelines

Meaningful only when the log type is `file`. Otherwise, it is ignored.

email-address

This command specifies the email address of a remote recipient of SMTP log messages.

Syntax

email-address *address*

Parameters

address
Specifies the remote email address.

Guidelines

The **email-address** command is only used when the log type is `smtp`.

Examples

Assign an email address.

```
# email-address techDesk@datapower.com
```

event

This command adds an event class and a priority to the log target.

Syntax

event *class priority*

Parameters

class Specifies the name of an event-class, which is a set of related events.

priority
Identifies the event priority.

Guidelines

Log priority is characterized (in descending order of importance) as `emerg`, `alert`, `critic`, `error`, `warn`, `notice`, `info`, and `debug`. The priority specifies that the appliance writes to the log all events that are greater than or equal to this criticality.

You can use the **show logging priority** command to display a list of event priorities.

You can use the **show logging event** command to display a list of event classes.

Example

Set which event classes and which event priorities to log.

```
# event schema error
# event xmlfilter error
# event crypto error
# event ssl error
# event auth warning
#
```

event-code

This command specifies an event code to add to the log target.

Syntax

event-code *code*

Parameters

code Identifies the hex value of the event code.

Guidelines

This command allows only messages that contain specified event codes to be written to the current log target. Thus, it is possible to create a log target that collects only messages for a particular set of event codes. For example, "Operational State down."

Use the **View List of Event Codes** from the GUI to view a list of all event codes.

Example

Create a file-based log target that contains only XML parser events.

```
# type file
# event-code 0x00030001
# event-code 0x00030002
# event-code 0x00030003
# event-code 0x00030004
# event-code 0x00030005
# event-code 0x00030006
# event-code 0x00030007
# event-code 0x00030008
# event-code 0x00030009
# event-code 0x0003000a
#
```

event-detection

This command suppresses identical events in the log target.

Syntax

event-detection { **on** | **off** }

Parameters

on Suppresses the writing of identical events to the log for the specified suppression period.

off Identical events are written to the log. This setting is the default value.

Guidelines

The **event-detection** command allows for the suppression of identical log events that are generated by the same configuration object over a configurable time period. When enabled, the log target retains a reference to each processed event for the duration set with the **suppression-period** command. Until this period expires, the log target does not process the same event from the same configuration.

event-filter

This command specifies an event code to exclude from the log target.

Syntax

event-filter *code*

Parameters

code Specifies the hex value of the event code.

Guidelines

Event filters provide for the exclusion of log messages that contain specified event codes from the log target.

Use the **View List of Event Codes** from the GUI to view a list of all event codes.

Example

Create a file-based log excluding XML parser events.

```
# type file
# event-filter 0x00030001
# event-filter 0x00030002
# event-filter 0x00030003
# event-filter 0x00030004
# event-filter 0x00030005
# event-filter 0x00030006
# event-filter 0x00030007
# event-filter 0x00030008
# event-filter 0x00030009
# event-filter 0x0003000a
#
```

facility

This command specifies the syslog facility.

Syntax

facility *facility*

Parameters

facility
Identifies the syslog facility.

Guidelines

The **facility** command specifies the syslog facility. This command is meaningful only with syslog-based log targets.

Examples

Set to the `local0` syslog facility.

```
# type syslog
# local address 10.10.13.4
# remote-address 172.16.100.1
# facility local0
#
```

feedback-detection

This command suppresses events from the logging subsystem itself.

Syntax

feedback-detection { **on** | **off** }

Parameters

on Suppresses all log events that are triggered by the logging subsystem.

off Suppresses log events that are triggered by the target itself, but writes events that are generated by other log targets. This setting is the default value.

Guidelines

The **feedback-detection** command allows for the suppression of log events that are triggered by the logging subsystem itself. Log targets always suppress log events that are triggered by the target itself but write events that are generated by other log targets. Under certain circumstances, with two or more log targets, these events can create a positive feedback loop that might cause resource contention. When enabled, feedback detection suppresses all log events that are triggered by the logging subsystem.

format

This command specifies the format in which events are added to the log.

Syntax

format *format*

Parameters

format

Indicates the format of log messages. The default value is `xml`.

cbe Specifies the log format follows the IBM Common Base Event specification.

csv Specifies the log format as comma-separated.

raw Specifies the log format as unformatted text.

text Specifies the log format as formatted text.

xml Specifies the log format as XML.

Guidelines

Use the **show logging format** command to list the available formats.

idle-timeout

This command sets the idle timer for the logging target.

Syntax

idle-timeout *seconds*

Parameters

seconds
Sets the number of seconds to wait before the appliance closes an inactive connection. Enter a value in the range 1 - 600. The default value is 15.

Guidelines

The **idle-timeout** command sets the number of seconds to wait before the appliance closes an established but inactive connection to the server.

Attention: If multiple log targets have the following configuration, they might share connections.

- The same local address and port.
- The same remote address and port.

Because of potential connection-sharing, set the same values for these log targets.

You can use the **active-timeout** and **connect-timeout** connection commands to modify other timeout values.

- The active timeout sets the number of seconds to wait before the appliance closes an established and active connection to the server.
- The connect timeout sets the number of seconds to wait to establish a connection to the server before an error message is generated.

ip-filter

This command adds an IP address to include in the log target.

Syntax

ip-filter *address*

Parameters

address
Identifies the network IP address.

Guidelines

The **ip-filter** command adds an IP address to include in the log target. Only log messages that are from the specified IP address are written to the log target. With this command, you can create a log target that collects log messages from only specific clients.

local-address

This command specifies the local address over which log events are transmitted to a remote recipient.

Syntax

local-address *address*

Parameters

address
Specifies the IP address of the interface.

Guidelines

The **local-address** specifies the local address over which syslog log events are transmitted to a remote recipient.

- For an SMTP log target, the **local-address** command is required. The log target uses TCP port 25.
- For a syslog-based log target, the **local-address** command is optional.
  - For syslog via UDP, the log target uses port 25.
  - For syslog via TCP, the log target uses port 514.

For all other log types, the **local-address** command is not used.

Example

Specify the local interface used to transmit log messages to an SMTP server.

```
# type smtp
# local-address 10.10.13.4
```

local-file

This command specifies a local file for log messages.

Syntax

local-file *URL*

Parameters

URL Specifies the file to store log messages and takes the `logstore:///`*`file`* form.

Guidelines

When the log type is `file`, the **local-file** command is required. For all other log types, it is not used.

local-ident

This command sets the string that the remote recipient uses to identify this log target.

Syntax

local-ident *identifier*

Parameters

identifier
Specifies the identifier for the log target.

Guidelines

The **local-ident** sets the string that a remote recipient uses to identify the log target. For an SMTP or syslog-based log target, this command is optional. For all other types of log targets, the value set by this command is ignored.

object

This command adds an object filter to the current log target.

Syntax

object *type* [*name*]

Parameters

type Specifies an object type. This filter restricts messages to only messages for that object type.

name Optional: Specifies the name of an instance of the selected object type.

Guidelines

Use the **object** command to enable a finer granularity of log content.

Object filters allow only those log messages that are generated by specific objects to be written to the log target. Object filters are based on object type and based on specific instances of that object type. Therefore, you can create a log target to collect log messages for an instance of a particular object type. For example, you can create a log target to write messages for the xyz service. Omit the instance name to filter on all instances of the specified type.

Examples

- Add an object filter to log messages for the domain-3 domain.

```
# object domain domain-3
#
```

- Add an object filter to log messages for the Gateway-1 Multi-Protocol Gateway.

```
# object XSLProxy Proxy-1
#
```

priority

This command sets the service-level priority for the log target.

Syntax

priority { **low** | **normal** | **high** }

Parameters

low Sets the DataPower service to receive below normal priority for scheduling and resource allocation.

normal Sets the DataPower service to receive normal priority for scheduling and resource allocation. This setting is the default value.

high Sets the DataPower service to receive above normal priority for scheduling and resource allocation.

Guidelines

The **priority** command sets the priority for log target event-flushing. When resources are in high demand, setting the priority to `high` can increase processing capacity of the log target but might have a negative effect on the throughput of the appliance.

rate-limit

This command sets the maximum rate of events to log.

Syntax

rate-limit *events/seconds*

Parameters

events/seconds
Sets the maximum number of transactions per second. Enter a value in the range 1 - 1000. The default value is 100.

Guidelines

The **rate-limit** command sets the number of log transactions per second.

- Remote log targets might receive more than this number of events within a second, depending on network latency and buffering. Because only a single TCP connection is made to the syslog server, a syslog over TCP log target is exclusive.
- For syslog over TCP log targets, the rate limit is the maximum number of events that are transmitted over the connection within 1 second. A value of 0 disables rate-limiting by the log target.

This command is meaningful for an NFS, SMTP, SOAP, or syslog-based log target. Otherwise, it is ignored.

Example

Limit transactions to a maximum of 50 per second.

```
# rate-limit 50
#
```

remote-address

This command specifies the destination address of log messages or the log itself.

Syntax

remote-address *host*

Parameters

host Identifies the host name or IP address of the remote server.

Guidelines

The **remote-address** command specifies the destination of transmitted log messages or the location of an uploaded log file. This command is relevant in the following situations.

- For an SMTP or syslog-based log target, as specified by the **type** command.
- For a file-based log target with an archive mode, as specified by the **archive-mode** command set to `upload`.

Use the **remote-address** command with the **remote-port** command to define the destination of transmitted log messages.

With TCP-based, network log targets, instead of specifying the IP address or host name of a remote server, you can specify the name of a load balancer group. In this situation, the same load balancer group must be assigned to the `default` XML manager in the domain with the XML Manager **loadbalancer-group** command. To create a load balancer group, use the Global **loadbalancer-group** command.

To establish a secure connection to a remote server, use the **remote-server** and **remote-port** commands to set the values to that of a local SSL Proxy service on the appliance. The local SSL Proxy service, as defined by the Global **sslforwarder** command, can then forward log messages over a secure connection to the remote server.

Examples

- Specify the address of an SMTP server. Uses the default TCP port of 25.

```
# type smtp
# local address 10.10.13.4
# remote-address ragnarok.datapower.com
#
```

- Specify the address of a syslog daemon. Uses the default UDP port of 514.

```
# type syslog
# local address 10.10.13.4
# remote-address 172.16.100.1
#
```

- Specify the recipient address for an uploaded log file.

```
# type file
# archive-mode upload
# remote-address 172.16.100.1
#
```

remote-directory

This command specifies the remote directory where uploaded logs are stored.

Syntax

remote-directory *file-path*

Parameters

file-path

Identifies the writable remote directory that stores uploaded log files.

Guidelines

The **remote-directory** command is used only in the following situations:

- The log type is `file`.
- The archive mode is `upload`.
- The upload method is `scp`, `ftp`, or `sftp`.

Note: The `scp` upload method is deprecated.

To denote an absolute directory from the root directory, specify a single forward slash character or equivalent encoded character (`%2F`) before the fully qualified file name.

- For SCP or SFTP, specify */file-path*.
- For FTP, specify `%2F`*file-path*.

The path in the URL resolves to `//`*file-path* for SCP or SFTP and `/%2F`*file-path* for FTP.

To denote a directory that is relative to the user's home directory, do not specify a forward slash character or equivalent encoded character before the fully qualified file name. For example, specify *file-path*. The path in the URL resolves to */file-path*.

Examples

- Specify the remote directory for an uploaded log file that is relative to the user's home directory.

```
# type file
# archive-mode upload
# upload-method sftp
# remote-address 172.16.100.1
# remote-port 2121
# remote-directory logs/
#
```

- Specify the remote directory for an uploaded log file that is absolute to the root directory.

```
# type file
# archive-mode upload
# upload-method ftp
# remote-address 172.16.300.254
# remote-port 2123
# remote-directory %2Flogs/
#
```

remote-login

This command specifies the user name to use to upload a log file to a remote server.

Syntax

remote-login *user* [*password*]

Parameters

user Specifies the user name to access a recipient of uploaded logs.

password
Specifies the password for the user name.

Guidelines

The **remote-login** command is used only if the log type is `file` and the archive-mode is `upload`.

Without a password, you must specify it during the upload session.

Examples

Specify the recipient address, user name, password, and remote directory for an uploaded log file.

```
# type file
# remote-address 172.16.100.1
# remote-login jrb brj
# remote-directory logs/
#
```

remote-port

This command specifies the listening port on the remote server.

Syntax

remote-port *port*

Parameters

port Specifies the destination port that monitors traffic. The default value is 514.

Guidelines

The **remote-port** command specifies the listening port on the remote server. This command is relevant only when for an SMTP or syslog-based log target, as specified by the **type** command.

Use the **remote-port** command with the **remote-address** command to define the destination of transmitted log messages.

You can use SSL to establish a secure connection to a remote server. For this configuration, set the values of the **remote-server** and the **remote-port** commands to the values of a local SSL Proxy on the appliance. The local SSL Proxy, as defined by the Global **sslforwarder** command, can then forward log messages over a secure connection to the remote server.

Example

Set the address of an SMTP server that listens on port 5400.

```
# type smtp
# local address 10.10.13.4
# remote-address ragnarok.datapower.com
# remote-port 5400
#
```

rotate

This command sets the maximum number of file rotations.

Syntax

rotate *count*

Parameters

count Specifies how many times to rotate a log file. Enter a value in the range 1 - 100. The default value is 3.

Guidelines

The **rotate** command specifies the maximum number of rotations for the log file.

Depending on the appliance type, the location of the file can be the local file system or the hard disk array.

Assuming a file name of `CryptoLog` and three rotations, the directory that contains the log file can contain the following local files.

CryptoLog
The current log file.

CryptoLog1
The log file that was most recently archived.

CryptoLog2
The log file that was next most recently archived.

CryptoLog3
The oldest log file.

This command is meaningful only when one of the following conditions is met.

- The log type for the **type** command is `file`, and the archival mode for the **archive-mode** command is `rotate`.
- The log type for the **type** command is `nfs`.

sender-address

This command specifies the email address of the sender.

Syntax

sender-address *address*

Parameters

address
Specifies the local email address.

Guidelines

The **sender-address** command is only used when the log type is `smtp`.

size

This command sets the maximum size of a local log file.

Syntax

size *KB*

Parameters

KB Specifies the maximum size of the file in KB. Enter a value in the range 100 - 50000. The default value is 500.

Guidelines

The **size** command sets the maximum size of a local log file in KB.

Depending on the appliance type, the location of the file can be the local file system or the hard disk array.

This command is only meaningful when the log type, as specified by the **type** command, is `file`.

smtp-domain

This command specifies the domain name of the SMTP client.

Syntax

smtp-domain *domain*

Parameters

domain
Specifies the fully qualified domain name of the SMTP client.

Guidelines

The **smtp-domain** command specifies the fully qualified domain name of the SMTP client. This domain name is sent as part of the SMTP session invitation (**HELO** command).

This command is meaningful only in the following situations:

- The log type, as specified by the **type** command, is `smtp`.
- The log type, as specified by the **type** command, is `file`, and the upload method, as specified by the **upload-method** command is `smtp`.

Example

Set the recipient of SMTP domain.

```
# type smtp
# smtp-domain popServer-1.example.com
#
```

soap-version

This command specifies the version of SOAP to use.

Syntax

soap-version { **soap11** | **soap12** }

Parameters

soap11 SOAP targets use SOAP 1.1. This setting is the default value.

soap12 SOAP targets use SOAP 1.2.

Guidelines

When the log type is `soap`, specifies the version of SOAP for use by SOAP log targets.

ssl

This command associates an SSL proxy profile for SOAP-based log over HTTPS.

Syntax

ssl *name*

no ssl

Parameters

name Specifies the name of an SSL proxy profile.

Guidelines

Use the **ssl** command associates an SSL proxy profile for SOAP-based log over HTTPS. This command is meaningful only when the target type is SOAP and URL uses HTTPS.

To create an SSL proxy profile, use the Global **sslproxy** command.

To remove the association of the SSL proxy profile, use the **no ssl** command.

suppression-period

This command sets the interval to suppress identical events.

Syntax

suppression-period *seconds*

Parameters

seconds
Specifies the interval in seconds to suppress identical events. The default value is 10.

timestamp

This command specifies the time stamp format.

Syntax

timestamp { **numeric** | **syslog** }

Parameters

numeric
Specifies a numeric time stamp format. This setting is the default value.

syslog Specifies a syslog time stamp format.

trigger

This command creates an event trigger.

Syntax

trigger *message-ID* [*expression*] *only-once only-this-trigger command*

no trigger *message-ID* [*expression*] *only-once only-this-trigger command*

Parameters

message-ID
Specifies the message identifier that will, when logged, trigger the command.

expression
Optionally specifies a regular expression that must match the body of the message to trigger the command.

only-once
Indicates the behavior when the trigger criteria are met more than one time.

on This command is triggered only the first time that the trigger criteria are met.

off This command is triggered each time that the trigger criteria are met. This setting is the default value.

only-this–trigger
Indicates how other event triggers behave that have the same trigger criteria.

on This command is triggered, but other commands that would be triggered by the same message ID are not. This setting is the default value.

off Subsequent event triggers, that share this message ID, are also triggered.

command
The command that runs when the trigger criteria are met. Commands must be separated by semicolons. If the command contains a space, enclose the value in double quotation marks.

Guidelines

Use the event trigger to run specific commands when specific messages appear in the system logs.

Examples

- Start a packet capture when the specified message is logged. The *only-once* parameter is set to `on` so that multiple packet captures are not initiated. The *only-this-trigger* parameter is set to `on` so that the stop packet capture command is not triggered immediately.

```
# trigger "0x99999" "" on on "interface eth0; packet-capture
    temporary:///capture -1 250"
#
```

- Stop the packet capture.

```
# trigger "0x99999" "" on on "interface eth0; no packet-capture
    temporary:///capture"
#
```

type

This command identifies the logging model.

Syntax

type *model*

Parameters

model Sets the logging model.

cache Writes log entries to memory.

console
Writes log entries to the console screen.

file Writes log entries to a file on the appliance.

nfs Writes log entries to a file on a remote NFS mount.

smtp Forwards log entries as email messages to a specified recipient.

soap Forwards log entries as SOAP messages to a specified recipient.

syslog Forwards log entries to a remote syslog daemon over UDP.

syslog-ng
Deprecated. Use `syslog-tcp`.

syslog-tcp
Forwards log entries to a remote syslog daemon over TCP.

Guidelines

For all log types, use the **event** command to specify log contents.

- For `cache`, requires no configuration beyond the identification of log type. You can use the **format**, **size**, and **timestamp** commands to customize log behavior.
- For `console`, no additional configuration is required. You can use the **format** and **timestamp** commands to customize log behavior.
- For `file`, you must specify the name of the log file with the **local-file** command. You can use the following commands to customize log behavior.
  - The **size** command to specify the file size.
  - The **rotate** command to specify the number of file rotations.
  - The **backup** command to specify a backup log.
  - The **format** command to specify the log format,
  - The **timestamp** command to specify the time stamp method.
  - The **archive-mode**, **upload-method**, **local-ident**, **remote-address**, **email-address**, and **sender-address** commands to specify an archival and retrieval method.
- For `nfs`, no **network-file** event messages for network file access are logged to NFS log targets. To troubleshoot `network-file` events, use a log target type other than NFS.
- For `smtp`, you must use the **local-address** and **remote-address** commands to specify source and destination IP addresses and the **email-address** command to specify a destination email address. You can use the following commands to customize log behavior.
  - The **rate-limit** command to control event message flow.
  - The **local-port** and **remote-port** commands to identify nonstandard source or destination ports.
  - The **sender-address** command to specify a pseudo recipient email address.
  - The **local-ident** command to specify a local identifier.
  - The **backup** command to specify a backup log.
  - The **format** command to specify the log format.

- For `soap`, you must use the **`url`** command to specify the recipient of SOAP-enveloped log events and use the **`rate-limit`** command to throttle the event flow. You can use the **`backup`** command to specify a backup log.
- For syslog-based logs, you must use the **`remote-address`** command to specify a remote IP address. You can use the following commands to customize log behavior.
  - The **`remote-port`** command to specify a nonstandard destination port number.
  - The **`local-address`** and **`local-port`** commands to specify a specific local interface for syslog transmissions.
  - The **`local-ident`** command to specify a local identifier.
  - The **`facility`** command to identify the syslog facility.

See the Global **`sslforwarder`** command for information on enabling the transmission of syslog events over a secure TCP connection.

upload-method

This command sets the protocol to upload a file-based log to a remote site.

Syntax

`upload-method` *protocol*

Parameters

protocol

Specifies the protocol to upload the report. The default value is `scp`, which is deprecated.

`ftp` Identifies the File Transfer Protocol.

`scp` Deprecated - Identifies the Secure Copy Protocol (SCP).

`sftp` Identifies the Secure File Transfer Protocol.

`smtp` Identifies the Simple Mail Transfer Protocol.

Guidelines

The **`upload-method`** command sets the protocol to upload a file-based log to a remote site. This command is used in the following situations:

- The log type that is set by the **`type`** command is `file`.
- The archive-mode that is set by the **`archive-mode`** command is `upload`.

The SCP upload method is deprecated. If you use SCP to upload multiple log targets at the same time, a system error might occur. The appliance supports only one SCP connection at a time. To minimize the risk when you use SCP, configure log targets with different settings, such as different event subscriptions and different log sizes.

Examples

Provide the required information (transfer protocol, recipient address, user name, password, and remote directory) to upload a file-based log to a remote storage site.

```
# type file
# upload-method sftp
# remote-address 172.16.100.1
# remote-login jrb brj
# remote-directory logs/
#
```

url

This command sets the destination for SOAP-based log entries.

Syntax

url *URL*

Parameters

URL Identifies the destination.

Guidelines

The **url** command sets the destination for SOAP-based log entries. This command is used only if the log type is `soap`.

Examples

Set the recipient of SOAP log messages.

```
# type soap
# url http;//ragnarok.example.com/logs
#
```

Monitoring and reporting commands

You use the monitoring and reporting commands to view the status of queue manager and of the appliance.

Use the status command to view general information about the appliance and specific information about a queue manager.

Use the show command to view details of the status of the appliance.

status

Reports disk usage, CPU usage, and memory usage across the appliance or for a specific queue manager. Also reports additional information for a queue manager running in a high availability configuration, or a disaster recovery configuration.

Purpose

You can use the **status** command to get information about the disk usage, CPU usage, and memory usage for the appliance or for a specific queue manager.

If the queue manager is running in a high availability configuration, the following information is also reported:

- The high availability role of the queue manager (reported as `Primary` or `Secondary`).
- The current high availability status:

Normal
The appliances in the disaster recovery configuration are operating normally.

This appliance in standby mode
This status means that the appliance has been suspended (by using the **sethagrp -s** command).

Secondary appliance in standby mode
This status means that the other appliance in the HA pair has been suspended (by using the **sethagrp -s** command).

Both appliances in standby mode
This status means that both appliances in the HA pair have been suspended (by using the **sethagrp -s** command).

Secondary appliance unavailable
This status means that the connections to the other appliance in the HA pair have been lost.

Partitioned
If the two appliances lose the replication interface, the queue managers are immediately flagged as partitioned. The running queue manager might accumulate out-of-sync data. The other queue manager remains in standby with no out-of-sync data. When the connection is remade, the partitioned state is automatically resolved and replication is resumed.

Synchronization in progress
This status is displayed when the primary queue manager is replicating data to the secondary queue manager.

Inactive
The queue manager is inactive on both appliances in the HA pair.

- The preferred appliance setting for the queue manager, set to `This Appliance` or `Other Appliance`.
- The percentage complete of a synchronization operation. This information is shown only when the status is `Synchronization in progress`.
- The estimated time at which a synchronization will complete. This information is shown only when the status is `Synchronization in progress`.
- The amount of out-of-sync data that exists on this instance of the queue manager. This is the amount of data written to this instance of the queue manager since it entered the partitioned state. This information is shown only when the status is `Partitioned`.

If the queue manager is running in a disaster recovery configuration, the following information is also reported:

- The disaster recovery role of the queue manager (reported as `Primary` or `Secondary`).
- The current disaster recovery status:

Normal
The appliances in the disaster recovery configuration are operating normally.

Synchronization in progress
This status can mean that initial replication is completing, or there has been a failure of the disaster recovery replication network and the queue manager has switched into synchronization mode to catch up as quickly as possible.

Partitioned
Queue manager data on the appliances is out of step, and cannot be automatically resolved. The **makedrprimary** and **makedrsecondary** commands must be used to resolve the situation. When this status is displayed on one of the appliances in a disaster recovery pair, the other appliance might display the **remote appliance unavailable** status, because the connection was lost before it detected the partitioned status.

Remote appliance(s) unavailable
The status means that the connection to the other appliance in the disaster recovery configuration has been lost.

Inactive
The queue manager is in the secondary role on both appliances.

Inconsistent
This status is shown only when the queue manager is in the secondary role and synchronization is in progress. If you use the **makedrprimary** command on a queue manager that is in this state, the queue manager reverts to the snapshot of its data that was taken before it entered the inconsistent state.

Reverting to snapshot
This status is shown when the queue manager is in the secondary role, and the **makedrprimary** command is issued when the queue manager is in the `inconsistent` state. The queue manager is reverted to the current snapshot of its data such that it can run.

Remote appliance(s) not configured
This status is shown when the **crtdrprimary** command has been run, to specify that a queue manager has the primary role, but no **crtdrsecondary** command has been run on the other appliance in the disaster recovery pair.

- The percentage complete of a synchronization operation. This information is shown only when the status is `Synchronization in progress`.
- The estimated time at which a synchronization will complete. This information is shown only when the status is `Synchronization in progress`.
- The amount of out-of-sync data that exists on this instance of the queue manager. This is the amount of data written to this instance of the queue manager since it entered the partitioned state. This information is shown only when the status is `Partitioned`.
- The percentage complete of a reversion to snapshot operation. This information is shown only when the status is `Reverting to snapshot`.

If the queue manager is part of both an HA and a DR configuration, then both HA and DR information is displayed.

Syntax

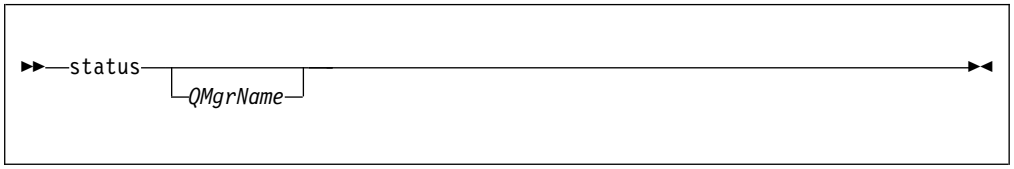

Parameters

QMgrName

Specifies the name of the queue manager for which the status summary is returned.

If this parameter is omitted, a summary of all disk and memory usage on the appliance is returned.

Usage Notes

- This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes mqa(mqcli)#. To enter the IBM MQ administration mode, enter `mqcli` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- The information that is returned for the appliance includes the following information:
  - The size and usage of the system memory
  - The CPU usage of the system
  - The size and usage of the internal disk. If the appliance has queue managers that belong to a disaster recovery configuration, the size information includes the space reserved for the snapshot logical volume for each disaster recovery queue manager.
  - The size and usage of the system volume
  - The number of FDCs and the disk space used
  - The disk space used by trace
- The information that is returned for a queue manager includes the following information:
  - The queue manager name
  - The queue manager status
  - The CPU usage of the queue manager
  - The memory usage of the queue manager. If this is a disaster recovery queue manager, this figure does not include the additional memory required for the snapshot image. Note that creating a primary queue manager in a disaster recovery configuration fails if there is insufficient memory bor both the queue manager data, and the snapshot of the queue manager data.
  - The amount of the queue manager file system used by the queue manager
- The information that is returned for a high availability queue manager can also include the following information:
  - The operational state of the HA group
  - The replication status of the queue manager (if synchronization is in progress)
  - The preferred appliance for the queue manager
  - Whether a partitioned situation has been detected, and if it has, the amount of 'out-of-sync' data held
- The information that is returned for a disaster recovery queue manager can also include the following information:
  - The disaster recovery role (primary or secondary)
  - The disaster recovery status
  - The percentage complete if synchronization is in progress
  - The estimated time to completion if synchronization is in progress
  - The amount of out-of-sync data if the disaster recovery system is partitioned

- The percentage complete if reversion to snapshot is in progress
- The number of logical writes not yet completed by the primary instance of a queue manager to the secondary instance.
- The number of logical bytes not yet written by the primary instance of a queue manager to the secondary instance.

Examples

- The following command returns a report for the appliance:
  ```
  status
  ```
- The following command returns a report for a specific queue manager, `QM1`:
  ```
  status QM1
  ```

show

Use the show command to view the status of the appliance or an aspect of the appliance configuration.

Type `show` on its own to see a list of arguments.

The command has the following format:

```
show status_provider
```

Where *status_provider* identifies the status that the command displays. The available status providers are described in the following sub topics.

or

```
show configuration_type
```

Where *configuration_type* can identify one of the appliance configuration types, as shown in the following table:

Table 37. Arguments for showing configuration with the show command

Argument	Description
aliases	Lists the currently defined command macros.
cert-monitor	Displays the current cert-monitor configuration.
certificate	Shows current certificate aliases.
dns	Displays the current DNS configuration, see "DNS commands" on page 490.
dns-cache	Shows the DNS hosts held in the DNS cache.
ethernet	Displays the current Ethernet configuration, see "Ethernet commands" on page 495.
features	
host-alias	Shows current host aliases.
image	Shows details of the current firmware image.
known-hosts	Shows current SSH known hosts. See "**known-host**" on page 516.
link-aggregation	Displays the current link-aggregation configuration, see "Link aggregation commands" on page 542.
name-servers	Shows currently defined name servers.
netarp	
network	Displays the current network configuration, see "Network commands" on page 599.

Table 37. Arguments for showing configuration with the show command (continued)

Argument	Description
ntp-service	Displays the current NTP service configuration, see "Date, time, and locale configuration commands" on page 478.
search-domains	Lists domain-suffixes in the search table for non-qualified domain names.
static-hosts	Lists currently defined static hosts.
syslog	Shows where logging messages are currently forwarded to.
throttle	Displays the current throttle configuration.
web-mgmt	Displays the current web management configuration, see "Web management service commands" on page 623.

`show authcookie-cache-status:`

This command lists information about AuthCookie cache entries.

Syntax

`show authcookie-cache-status`

Guidelines

The **`show authcookie-cache-status`** command lists the following information about AuthCookie cache entries:

- The hashed key for the cache entry
- The data value that the cache entry stores
- The time when the cache entry expires
- The owner of the cache entry
- The domain where the cache entry is created

The AuthCookie cache is used by the HTML forms login policy and the social login policy.

- The HTML forms login policy uses the AuthCookie cache to create and refresh session cookies.
- The social login policy uses the AuthCookie cache to store and retrieve OIDC client state information.

To remove the data from the cache, use the **`Crypto authcookie-cache-delete`** command.

`show clock:`

This command displays the current time and appliance uptime.

Syntax

`show clock`

Guidelines

The **`show clock`** command produces the same results as the **`show time`** command.

Output

```
# show clock

    Local Time: Fri Nov 30 12:03:02 2012
     Time Zone: EST
Time Zone Spec: EST5EDT,M3.2.0/2:00,M11.1.0/2:00
Uptime: reload: 0 days 00:00:17
Uptime: reboot: 35 days 11:23:04
```

show cpu:

This command displays average CPU utilization data for the last 10 seconds, 1 minute, 10 minutes, 1 hour, and 1 day.

Syntax

show cpu

show ethernet-counters:

This command displays the statistics about Ethernet devices.

Syntax

show ethernet-counters

Availability

Physical appliances only.

Guidelines

The **show ethernet-counters** command provides the detailed statistics available for each Ethernet device on the system. Different interfaces on different appliance models keep different statistics, so some entries might not have values.

show ethernet-mau:

This command displays the device-specific status of all Ethernet MAUs on the system.

Syntax

show ethernet-mau

Availability

Physical appliances only.

Guidelines

The **show ethernet-mau** command displays the device-specific status of all Ethernet MAUs on the system.

show ethernet-mii-registers:

This command displays the MII registers of all Ethernet MAUs.

Syntax

show ethernet-mii-registers

Availability

Physical appliances only.

Guidelines

The **show ethernet-mii-registers** command displays the MII Registers of all Ethernet MAUs on the system with access to the MII registers.

show file:

If used in config mode this command displays a specified printable file. Otherwise displays total and free space details for the appliance.

Syntax

In config mode:
show file *URL*

Not in config mode:
show file

Parameters

URL Identifies the URL of the file to display. The URL takes the *directory*:///*file* format.

Where:

directory
Specifies a directory on the appliance.

file Specifies the name of a file in the directory.

Guidelines

You cannot use the **show file** command to display files in the cert: directory.

show firmware:

This command displays the current firmware version, with image type and installation date.

Syntax

show firmware

Guidelines

The **show firmware** command provides a subset of the details of the **show firmware-version** command. The **show firmware** command includes information about whether the current firmware image is the primary or secondary installation image and the date when the image was installed.

The command displays the following information.

- The type of the firmware installation type. Can be primary or secondary.
- The version of the firmware image.
- The specific build of the firmware image.
- The build date of the firmware image.
- The date that the firmware was installed.
- The number of times the firmware image was restarted. The count is from the initial firmware load on the appliance until the current time. The count is independent of firmware version.

show firmware-version:

This command displays the current firmware version, without image type and installation date.

Syntax

```
show firmware-version
```

Guidelines

The **show firmware-version** command provides information about the current firmware version. This command provides the same details as the **show version** command, but it does not provide the versions of the licenses that are available with the **show library-version** command.

The **show firmware-version** command does not include information about whether the current firmware image is the primary or secondary installation image or the date on which the image was installed. For these details, use the **show firmware** command.

The **show firmware-version** command displays the following information.

- The serial number of the appliance. This number is unique to each appliance.
- The version of the firmware.
- The build number of the firmware.
- The date on which the firmware was built.
- The build number of the system health monitor (watchdog).
- The version number of the installed DataPower appliance manager (DPOS).
- The version number of the running DataPower appliance manager.
- The type of XML accelerator.
- The IBM Machine Type of the appliance.
- The IBM Model Type of the appliance.

show ipaddress:

This command provides IP address information about interfaces.

Syntax

```
show ipaddress
```

Guidelines

The **show ipaddress** command provides IP address information about interfaces.

Output

The output includes the following data about each interface.

- The name of the interface.
- The primary address for the interface: IP version, address, and netmask.

show ipmi-lan-channel:

This command shows the current IPMI LAN channel configuration.

Syntax

show ipmi-lan-channel

show ipmi-user:

This command shows the current IPMI user configuration.

Syntax

show ipmi-user

show key:

This command lists the keys that have been configured.

Syntax

show key

Guidelines

The **show key** command lists keys.

Output

```
mqa# show key

key: iop-mgmt-key [down]
-----------------
 admin-state enabled
 file-name cert:///dtxkey.pem
```

show link-aggregation-member-status:

This command lists members in link aggregation interfaces.

Syntax

show link-aggregation-member-status

Guidelines

The **show link-aggregation-member-status** command lists the members in link aggregation interfaces. To view statistics for aggregate interfaces, use the **show link-aggregation-status** command. The output shows the following data about each aggregate interface.

- The aggregate interface's index.
- The aggregate interface's name: Kernel and configuration.
- The member interface's index.
- The member interface's name in the configuration.
- The ID that is assigned to the aggregator.

show link-aggregation-status:

This command provides statistics for aggregate interfaces.

Syntax

show link-aggregation-status

Guidelines

The **show link-aggregation-status** command provides statistics for aggregate interfaces. To view the members in a link aggregation, use the **show link-aggregation-members-status** command. The output shows the following data about each aggregate interface.

- The aggregate interface's index.
- The interface's name: Kernel and configuration.
- The maximum transmission unit (MTU) for the aggregation.
- The mode for link aggregation.

  Active-backup (ab)
  The mode is active-backup. One link is active, and the other link is backup.

  LACP The mode uses the Link Aggregation Control Protocol (LACP). The IEEE 802.1AX-2008 (formerly IEEE 802.3ad) standard defines LACP.

  Transmit load balancing (tib)
  The mode uses transmit-based load balancing.

  Unsupported
  The mode is not supported. Contact IBM Support.
- The link state of the aggregate interface. Values are `ok` or `no link`. If `no link`, the interface is not responding to the network.
- Which link, if any, is primary when the link aggregation mode is active-backup.
- The active link when the link aggregation mode is transmit-based load balancing or active-backup.
- With the LACP aggregation mode, the distribution algorithm for outbound packets among the active physical interfaces. The policy should match the configuration.

  layer2 Indicates that the hash function uses only MAC addresses, which is layer 2 data.

layer2+3
Indicates that the hash function uses both MAC addresses and IP addresses. MAC addresses are layer 2. IP addresses are layer 3. This setting is the default value.

layer3+4
Indicates that the hash function uses both IP addresses and TCP/UDP ports. IP addresses are layer 3. TCP/UDP ports are layer 4. This approach is not strictly compliant to the LACP standard.

- With the LACP aggregation mode, the transmission rate that members send LACP frames. The rate can be either `fast` or `slow`. In LACP, the frame uses the term *LACPDU*.
- With the LACP aggregation mode, the selection policy. The policy should match the configuration.

  bandwith
  Chooses the aggregator with the highest bandwidth.

  count Selects the aggregator with the most NICs.

  stable Does not change the aggregator when a better one is available.
- With the LACP aggregation mode, the active aggregator's ID.
- With the LACP aggregation mode, the number of available ports for the active aggregator.
- With the LACP aggregation mode, the value of the operational key that the Actor assigned to the port.
- With the LACP aggregation mode, the value of the operational key that the Partner assigned to the port.
- With the LACP aggregation mode, the Partner's physical MAC address.
- The MII link state of the aggregation. Values are `ok` or `no link`. If `no link`, the interface is not responding to the network.

show link:

This command provides status about all interfaces on the appliance.

Syntax

show link

Guidelines

The **show link** command provides status about all interfaces on the appliance. The output includes the following data about each interface.

- The name of the interface.
- The link state of the interface. Values are `ok` or `no link`. If `no link`, one of the following conditions exists.
  - The interface is not responding to the network.
  - The Ethernet interface is configured for link aggregation but is not assigned to an aggregate interface.
- The actual fixed, configured, or negotiated Ethernet `PHY` mode, including speed and duplex. All network interfaces, other than Ethernet interfaces, report a value of none.

- The type of interface: Ethernet, VLAN, aggregate, or other. The other interface type refers to network interfaces that you cannot configure. Examples of other interface types include `gre0`, `ip6tnl0`, `lo`, `sit0`, and `usb0`. Not all appliances have all of these other interface types.
- The maximum transmission unit (MTU), or largest packet, size that can be sent or received on the interface.
- The associated aggregate interface, if any.
- The physical address for the interface. The default value for Ethernet interfaces is programmed in the Ethernet hardware.

`show load:`

This command displays task level system usage.

Syntax

`show load`

Guidelines

The **`show load`** command displays the system usage by task. Use this command with the **`load-interval`** command to monitor system load.

The **`show load`** command displays memory usage for many tasks across the system, not for just the main task. The total memory usage from the **`show load`** command might be higher than the memory usage that displays with the **`show memory`** command. This discrepancy is because the **`show load`** command displays memory for these additional tasks.

The **`show load`** command displays the following status information.

- The instance of the task.
- The name of the task.
- The percentage of total load for the tasks during the measurement interval. A load above 90 indicates that the task is at or near load capacity. High load values are not necessarily indicative of a problem, if transaction latencies are not affected. Low load values mean that the task is looking for work. High load values mean that the task is doing work.
- The number of pending work items, including internal communication, in the queue for processing by the appliance. This number is for internal diagnostic tests and is not intended for capacity planning.
- The percentage of processor capacity that is used by this task.
- The percentage of memory that is used by this task. This ratio is relative to the amount of memory accessible to each task, rather than the memory accessible to the appliance as a whole. The cumulative task memory is not comparable to the memory usage as reported by the **`show memory`** command.
- The number of open file handles. The file count increases when the appliance accepts new connections or opens a file for reading. The count decreases when this file handle is closed. The file count value helps to identify which tasks are contributing to the high open file counts.

`show log:`

This command displays the appliance default log.

Syntax

show log

Guidelines

Use **show log** to display the default log. Use **show logging** to display the default or other log files.

show logging:

This command displays a specified appliance log.

Syntax

show logging *log-name* [*PCRE*]

show logging archive

show logging category [*category*]

show logging event

show logging format

show logging priority

show logging status [*name*]

show logging target [*name*]

show logging timestamp

show logging type [*type*]

show logging upload

Parameters

***name* [*PCRE*]**
Specifies the name of a log, and optionally displays only the events from the specified log that match the specified expression.

archive
Displays a list of available archival methods.

category [*category*]
Displays summary information about all active log categories, or displays summary information about the specified log category.

event Displays a list of supported event classes.

format Displays a list of supported log formats.

priority
Displays a list of event priorities.

status [*name*]
Displays summary information about all active log targets or a specific log target.

target [*target-name*]
Displays configuration data for all log targets or a specific log target.

timestamp
Displays a list of time stamp formats.

type [*type*]
Displays summary information about all available logging types, or displays detailed information about the specified log type.

upload Displays a list of available upload methods.

Guidelines

Use **show log** to display the default log.

show loglevel:

This command displays the log-level for logging targets.

Syntax

show loglevel

Guidelines

The **show loglevel** command shows the minimum log level for logging targets.

Log messages are characterized in descending order of criticality as `emerg`, `alert`, `critic`, `error`, `warn`, `notice`, and `info`.

Example

Show the log level of logging targets in the current application domain. The output shows that the `default-log` logging target in the `default` domain is set to `error`.

```
# show loglevel
Minimum log level for 'default-log' is 3 (error) in 'default' domain
```

show ndcache:

This command shows the status of IPv6 neighbor discovery translations on all interfaces.

Syntax

show ndcache

Guidelines

The **show ndcache** command shows neighbor discovery (ND) translations on all local interfaces. The output shows only complete cache entries. The output shows the following data about the network interface node:

- The primary address for the interface: IP version, address, and netmask.
- The physical (MAC) address of the interface.
- The name of the interface.

- The type of interface: Ethernet, VLAN, aggregate, or other. The other interface type refers to network interfaces that you cannot configure. Examples of other interface types include `gre0`, `ip6tnl0`, `lo`, `sit0`, and `usb0`. Not all appliances have all of these other interface types.
- The lifecycle state for the entry.

  `incomplete`
  Resolving neighbor and inaccessible.

  `reachable`
  Neighbor is valid and accessible.

  `stale` Neighbor is valid, but potentially reachable. State to be checked at first transmission.

  `delay` Awaiting verification from stale neighbor and inaccessible.

  `failed` Failed to resolve and inaccessible.

  `noarp` No attempt to resolve and inaccessible.

  `permanent`
  Neighbor is valid forever and always accessible.

  `probe` Awaiting confirmation from neighbor and inaccessible.

`show network-interface:`

This command shows generic status of all network interfaces on the appliance.

Syntax

`show network-interface`

Guidelines

The **`show network-interface`** command shows the following information about all network interfaces on the appliance.

- The type of interface: Ethernet, VLAN, aggregate, or other. The other interface type refers to network interfaces that you cannot configure. Examples of other interface types include `gre0`, `ip6tnl0`, `lo`, `sit0`, and `usb0`. Not all appliances have all of these other interface types.
- The name of the interface.
- The configured goal operational state of the interface.
- The operational state of the interface. States are `Up`, `Down`, `Unknown`, `Dormant`, `Not present`, or `Lower Layer Down`.
- The primary address for the interface: IP version, address, and netmask.
- The physical (MAC) address of the interface.
- The maximum transmission unit (MTU), or largest packet, size that can be sent or received on the interface.
- Statistics about received transactions:
  - The amount of data, in bytes, that was received successfully on the interface, which includes MAC-framing.
  - The number of packets that were received successfully on the interface and were passed to the network layer.
  - The number of packets that were not received because of errors in the packet or in the hardware.

  - The number of received packets that were not in error but were not passed to the network layer because of resource constraints.
- Statistics about transmitted transaction:
  - The amount of data, in bytes, that was transmitted successfully on the interface, which includes MAC-framing.
  - The number of packets that were transmitted successfully on the interface.
  - The number of packets that were not transmitted because errors on the network or in the hardware.
  - The number of packets that were not transmitted because the network layer generated packets faster than the physical network can accept them.

show ntp-refresh:

This command lists the refresh status for the current NTP server.

Syntax

show ntp-refresh

Guidelines

The **show ntp-refresh** command lists the following information about the current NTP server.

- The IP address of the last NTP server that was contacted.
- The results of the contact.
- The time after the refresh.

show password-map:

This command lists the defined aliases that have password maps.

Syntax

show password-map

Guidelines

The **show password-map** command list the number of passwords maps and their aliases.

Context

Available in only Crypto configuration mode.

Output

```
(config-crypto)# show password-map

2 password-map aliases
  george
  fyvush
```

show profile:

This command lists cryptographic profiles that have been configured.

Syntax

show profile

Guidelines

The **show profile** command lists cryptographic profiles.

Context

Available in only Crypto configuration mode.

show raid-array:

This command displays the status of the RAID array.

Syntax

show raid-array

Availability

Physical appliances only.

Guidelines

The **show raid-array** command displays the status of the RAID array.

- The reference number of the RAID card. The value is always 1.
- The reference number of this array. Numbering starts with 1.
- The identifier of the logical driver of which this array is a part.
- The RAID level of this array configuration.
- The number of physical drives for this array.
- The normalized size of this array in megabytes. The value is rounded down to an even multiple so that you can swap drives of the same nominal size but might not be the same raw size.

show raid-battery-module:

This command displays the information about the battery backup unit of the RAID controller.

Syntax

show raid-battery-module

Availability

Physical appliances only.

Guidelines

The **show raid-battery-module** command displays information about the battery backup unit (BBU) of the RAID controller. The BBU protects against the loss of cached data in the event of power failure.

- The reference number of the RAID card. The value is always 1.
- The type of BBU.
- The serial number of the BBU.
- The name of the BBU.
- The status of the BBU.

  chargeActive
  The battery is charging.

  dischargeActive
  The battery is discharging.

  i2cErrorsDetected
  The battery has inter-integrated circuit (I2C) errors.

  learnCycleActive
  The battery is in the learning cycle.

  learnCycleFailed
  The learning cycle for the battery failed because of errors in the learning cycle.

  learnCycleRequested
  A request was submitted for a learning cycle.

  learnCycleTimeout
  The learning cycle for the battery timed out.

  normal The battery is in a normal state.

  packMissing
  The battery is unplugged.

  periodicLearnRequired
  A request was submitted for a learning cycle to run battery maintenance.

  remainingCapacityLow
  The remaining capacity of battery is low.

  replacePack
  The battery needs to be replaced.

  temperatureHigh
  The battery temperature is high.

  undefined
  The battery state is undefined.

  voltageLow
  The battery voltage is low.
- The actual voltage of the battery in millivolts.
- The current through the battery terminals in milliamperes.
- The temperature of the battery in degrees Celsius.
- The designed capacity of the battery in milliampere-hour.
- The designed voltage of the battery in millivolts.

show raid-logical-drive:

This command displays the status of the RAID logical drive.

Syntax

show raid-logical-drive

Availability

Physical appliances only.

Guidelines

The **show raid-logical-drive** command displays the status of the RAID logical drive.

- The reference number of the RAID card. The value is always 1.
- The reference number of the logical drive.
- The name of the logical drive. The value is always `raid0`.
- The RAID level of the logical drive.
- The number of physical drives in the logical drive.
- The state of the logical drive in the volume.
- The state of the logical drive initialization.
- The read policy in effect.
- The write policy in effect.
- The cache policy in effect.
- The access policy in effect.
- An indicator of whether there are bad blocks in the logical drive.
- The total size of the logical drive in megabytes.

show raid-physical-drive:

This command displays the status of the RAID physical drive.

Syntax

show raid-physical-drive

Availability

Physical appliances only.

Guidelines

The **show raid-physical-drive** command displays the status of the RAID physical drive.

- The reference number of the RAID card. The value is always 1.
- The reference number of the physical disk. Numbering starts with 1.
- The reference number to which this array joins.
- The identifier of the logical driver of which this array is a part.
- The name of the logical drive. The value is always `raid0`.
- The location of the disk in the appliance.
- The overall state of the disk.
- The progress of a `rebuild`, `copyback`, `patrol`, or `clear` operation against the disk in percents.

- The exact size of this array in megabytes.
- The normalized size of this array in megabytes. The value is rounded down to an even multiple so that you can swap drives of the same nominal size but might not be the same raw size.
- The type of interface.
- The speed of the SAS interface. The speed is negotiated between the RAID controller and the physical disk.
- The SAS address of the disk.
- The vendor identification string for the disk.
- The product identification string for the disk.
- The vendor-provided revision string for the disk.
- The vendor-specific identifier for the disk. Normally, the value is unique for each physical drive.

show raid-ssd:

This command displays the remaining expected lifetime of the solid state disks in the RAID array.

Syntax

show raid-ssd

Availability

Physical appliances only.

Guidelines

The **show raid-ssd** command displays the expected remaining life of the SSDs. The lifetime of SSDs depends upon their workload (writes per day). The following information is dispayed:

- The disk number. The value is 1 or 2.
- The serial number of the SSD drive.
- The total GB written.
- The estimated life left of the SSD drive.

For example:

```
mqa# show raid-ssd

    Disk Number  SN        Total Written (Gib) Life Left (%)
    -----------  --------  ------------------- -------------
    1            ZAM11507  86                  100
    2            ZAM1150Q  86                  100
```

show route:

This command shows the routing table.

Syntax

show route

Guidelines

The **show route** command shows the routing table. This table describes the IP routes on the appliance. The table includes static and default routes from interface configurations and dynamic routes from discovery protocols.

For each route, the table provides the following information.

- The IP version, address, and netmask information for the destination.
- The type of interface: Ethernet, VLAN, or aggregate.
- The name of the interface.
- The IP version and address of the next hop gateway.
- The metric for the route. When the same destination has multiple routes, the route with the lowest metric is used.

show sensors-current:

This command displays the values for sensors that read electrical current.

Syntax

```
show sensors-current
```

Availability

Physical appliances only.

Guidelines

The **show sensors-current** command provides values for sensors that read electrical current. These sensors provide the current that certain components of the appliance use. This command returns an empty output on blades. The output for this command includes the following data.

- The name of the current sensor that is being monitored.
- The most recent reading of the current sensor in milliamperes. There are only three significant digits.
- The maximum allowable reading of the current sensor in milliamperes.
- Whether the current reading is OK or the reading exceeds the upper critical threshold. If the status is not OK, contact IBM support.

show sensors-fans:

This command displays the values for sensors that read the speed of the fans.

Syntax

```
show sensors-fans
```

Availability

Physical appliances only.

Guidelines

The **show sensors-fans** command provides values for sensors that read fan speed. This command returns an empty output on blades. The output for the **show sensors-fans** command includes the following data.

- The identifier for the fan.
- The fan speed in revolutions per minute (RPM).
- The lowest allowable reading of the fan speed sensor.
- The highest allowable reading of the fan speed sensor. If the maximum speed for a fan does not exist, no value is displayed.
- Whether the current fan speed is OK or the fan speed is below the lower critical threshold. If the status is not OK, contact IBM support.

show sensors-other:

This command displays the status of sensors that have true or false values.

Syntax

show sensors-other

Availability

Physical appliances only.

Guidelines

The **show sensors-other** command provides the state of sensors that have true or false (nonnumeric) values.

Which sensors the appliance provides depends on the appliance. The following list includes the sensors that different appliance provide, but the sensors are not limited to the ones in the list.

- Battery status
- Intrusion detection
- Hard disk drive status
- RAID controller status
- Power supply status, including failures, lack of input power, and so forth
- DRAM memory status
- PCI bus status
- CPU status

The output for the **show sensors-other** command includes the following data.

- The name of the sensor that is being monitored. The name is in the form of a descriptive predicate. If the value of the sensor is `true`, the descriptive predicate is `true`; if the value of the sensor is `false`, the descriptive predicate is `false`. For example, the sensor name `Intrusion detected` with a value of `false` means that intrusion is not detected.
- The value of the sensor. The value is either true or false.
- The status of the sensor. The status can be one of the following values.

  OK The reading of the sensor is normal.

Failure
The reading of the sensor indicates a problem or a failure in the appliance.

No reading
No reading of the sensor is available now. An internal hardware or software error might occur. Contact IBM Support.

Invalid
The DataPower software requested a sensor reading by using an invalid sensor identifier. An internal DataPower software error might occur. Contact IBM Support.

show sensors-temperature:

This command displays the values for sensors that read temperatures.

Syntax

```
show sensors-temperature
```

Availability

Physical appliances only.

Guidelines

The **show sensors-temperature** command provides values for sensors that read temperatures. These sensors provide the temperature of the air that flows through the appliance and key components of the appliance.

- Temperature of each internal CPU components

  For Type 8436 appliances, the sensors also read the temperature for each DIMM of the CPU components.
- Temperature of the internal inlet and outlet air temperatures

  For Type 8436 appliances, the System 1 sensor reads the inlet air temperature that is located in the front of the appliance. The System 2 sensor reads the outlet air temperature that is located in the rear of the appliance.

The output for the **show sensors-temperature** command includes the following data.

- The name of the temperature sensor.
- The most recent reading of the temperature sensor in degree Celsius. There are only three significant digits.
- The temperature at which a warning of high temperature occurs. If the temperature is above this value, investigate the cause and correct the problem.
- The temperature at which a critical error of high temperature occurs. If the temperature is above this value, correct the problem immediately.
- The temperature at which a risk of permanent damage to the appliance exists. If the temperature is above this threshold, correct the problem immediately. Otherwise, turn off the power supply of the appliance.
- Whether the current temperature is OK or exceeds the warning, critical, or danger threshold.

show sensors-voltage:

This command displays the values for sensors that read voltage.

Syntax

```
show sensors-voltage
```

Availability

Physical appliances only.

Guidelines

The **show sensors-voltage** command provides values for sensors that read voltages. These sensors provide the voltage of the power supplies and for other components of the appliance.

The output for the **show sensors-voltage** command includes the following data.

- The name of the voltage sensor.
- The most recent reading of the voltage sensor in millivolts. There are only three significant digits.
- The lowest allowable reading of the voltage sensor.
- The highest allowable reading of the voltage sensor.
- Whether the current voltage is OK or the voltage exceeds the upper or lower critical threshold. If the status is not OK, contact IBM support.

show services:

This command lists all local services that are listening for incoming connections.

Syntax

```
show services
```

Guidelines

Use the **show services** command to list all local services that are listening for incoming connections. For each entry in the table, the service that created the listener is shown.

An IP address of 0.0.0.0 indicates that the service is active on all interfaces.

Output

```
# show services

local IP        local port type      name
--------        ---------- ----      ----
0.0.0.0         5550       xml-mgmt  xml-mgmt
0.0.0.0         9090       web-mgmt  web-mgmt
```

show services-memory:

This command displays memory usage for active services.

Syntax

show services-memory

Guidelines

Use the **show services-memory** command to display a list of all active services and their memory usage in MB.

The output shows the memory usage in MB for each service within the following time frames.

- When the report is generated (current)
- During the last 60 seconds (1 minute)
- From the end of the first minute to the end of the fifth minute (1 - 5 minutes)
- From the end of the fifth minute to the end of the tenth minute (5 - 10 minutes)
- From the end of the tenth minute to the end of the first hour (10 minutes - 1 hour)
- From the end of the first hour to the end of the twelfth hour (1 - 12 hours)
- From the end of the twelfth hour to the end of the first day (12 hours - 1 day)
- From the end of the first day (lifetime)

show system:

This command displays the System Settings configuration.

Syntax

show system

show tcp-connections:

This command lists the number of TCP connections in specific states.

Syntax

show tcp-connections

Guidelines

The **tcp-connections** command list the number of TCP connections in the following states.

- established
- syn-sent
- syn-received
- fin-wait-1
- fin-wait-2
- time-wait
- closed
- close-wait
- last-ack
- listen
- closing

To list the current TCP connections, use the **show tcp-table** command.

show tcp-table:

This command lists the current TCP connections.

Syntax

show tcp-table

Guidelines

The **show tcp-table** command lists the current TCP connections. To list the number of TCP connections in specific states, use the **show tcp-connections** command.

show tcp:

This command lists the current TCP connections followed by the number of connections in each state.

Syntax

show tcp

Guidelines

The **show tcp** command lists the current TCP connections followed by the number of connections in each state. This command provides the details that are provided by the **show tcp-connections** and **show tcp-table** commands.

show throughput:

This command displays interface-specific traffic statistics.

Syntax

show throughput

Guidelines

The **show throughput** command displays interface-specific traffic statistics. The output includes the following tables.

- Received transactions in kilobits per second.
- Received transactions in packets per second.
- Transmitted transactions in kilobits per second.
- Transmitted transactions in packets per second.

Each table lists the count for each interface for the last 10 seconds, 1 minute, 10 minutes, 1 hour, and 24 hours.

Context

To view statistics, data collection must be active. Use the global **statistics** command to control data collection. If disabled, the command returns the `Statistics disabled` message.

show time:

This command displays the current time and appliance uptime.

Syntax

show time

Guidelines

The **show time** command produces the same results as the **show clock** command.

Output

```
# show time

    Local Time: Fri Nov 30 12:03:02 2012
     Time Zone: EST
Time Zone Spec: EST5EDT,M3.2.0/2:00,M11.1.0/2:00
Uptime: reload: 0 days 00:00:17
Uptime: reboot: 35 days 11:23:04
```

show users:

This command lists all users who are currently logged in to the appliance.

Syntax

show users

Output

```
# show users

 Session ID  Name  Connection    IP address   Login   Domain
 ----------  ----  ----------    ----------   -----   ------
```

show version:

This command displays the version of the firmware and libraries.

Syntax

show version

Guidelines

The **show version** command provides the combined details of the **show firmware-version** and **show library-version** commands.

Network commands

You can use the network commands to modify network settings on the IBM MQ Appliance.

The network commands can be run from the command line interface in network configuration mode. To enter network configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:

```
config
```

2. From global configuration mode, enter network settings mode:
   ```
   network
   ```
3. Type `exit` to leave the network settings mode and save your changes, then type `exit` again to leave global configuration mode.

arp-interval

This command sets the time interval between ARP attempts.

Syntax

arp-interval *milliseconds*

Parameters

milliseconds
Sets the time interval in milliseconds between ARP attempts. Enter a value in the range 500 - 5000. The default value is 500.

Guidelines

The **arp-interval** command sets the time interval between ARP attempts. The appliance waits the defined time before it tries a failed ARP request again.

Example

Set the interval to 100 milliseconds.

```
# arp-interval 100
```

arp-retries

This command sets the number of times the appliance attempts a failed ARP request.

Syntax

arp-retries *attempts*

Parameters

attempts
Sets the number of attempts for a failed ARP request. Enter a value in the range 1 - 64. The default value is 8.

Guidelines

The **arp-retries** command sets the number of times the appliance attempts a failed ARP request.

Example

Set the number of attempts to 5.

```
# arp-retries 5
```

block-traffic

This command indicates whether to block nonmanagement traffic for an invalid interface configuration.

Syntax

block-traffic { **on** | **off**

Parameters

on The appliance blocks nonmanagement traffic. The setting is the default value.

off The appliance allows all network traffic.

Guidelines

When at least one network interface has an invalid configuration, the **block-traffic** command indicates whether to block nonmanagement traffic or allow all network traffic.

For a production appliance, always enable this property. With this setting, the appliance supports only management traffic over Telnet, SSH, web management interfaces (WebGUI and Blueprint Console), and the XML management interface. Until you correct the problem, the appliance cannot accept and process client requests.

Example

For a test appliance, allow all network traffic.

```
# block-traffic off
```

destination-routing

This command controls how the appliance determines the route to return the response.

Syntax

destination-routing { **on** | **off** }

Parameters

on Interface selection is based on the best path to the client, irrespective of the service or receiving interface.

off Interface selection is based on the interface that is bound to the address of the service that generated the response. This setting is the default value.

Guidelines

The **destination-routing** command controls how the appliance determines the route to return the response to the originating client. The route to return the response is the outbound packet. The originating client is the destination of the outbound packet.

- When enabled, interface selection is based on the best path to the client, irrespective of the service or receiving interface. The best path is determined by static routes that are bound to the available interfaces.

  Note: Destination-based routing is for compatibility with an earlier version only. Enable destination-based routing only if an upgrade disables existing connectivity.

- When disabled, the default, interface selection is based on the interface that is bound to the address of the service that generated the response.
  - If the service is bound to a single address, responses are routed through the interface that is assigned to that address.
  - If the service is bound to multiple addresses, responses are routed through the receiving interface instead of the interface for the service that generated the response.

Example

Ensure that outbound packets originate from an interface that is bound to an address of the service that created the packet.

```
# destination-routing off
```

disable-interface-isolation

This command controls interface isolation.

Syntax

disable-interface-isolation { **on** | **off** }

Parameters

on Enables interface-isolation.

off Disables interface-isolation. This setting is the default value.

Guidelines

The **disable-interface-isolation** command controls whether to allow interface isolation. By default the appliance refuses to accept a packet on an interface other than the one bound to the destination address of the packet. As a security policy, the interface that receives a packet must be configured with the IP address that is the destination address of the packet. Enabling interface isolation relaxes the restriction.

Example

Allow any interface on the same subnet to accept a packet.

```
# disable-interface-isolation on
```

ecn-disable

This command controls ECN-capable TCP sessions.

Syntax

ecn-disable { **on** | **off** }

Parameters

on Stops the generation of ECN-capable TCP sessions.

off Generates ECN-capable TCP sessions. This setting is the default value.

Example

Stop the appliance from generating ECN-enabled TCP sessions.

```
# ecn-disable on
```

ephemeral-port-range

This command sets the starting port for the ephemeral port range.

Syntax

ephemeral-port-range *port*

Parameters

port Specifies the starting port of the ephemeral port range. By default, the appliance can use ephemeral ports in the range 10000 - 61000. You can override the default by defining a subset of ephemeral ports to not be ephemeral ports. These ports are ephemeral ports in the range 10000 - 32768. Even if you override, the appliance always has control of the ephemeral ports in the range 32769 - 61000.

Guidelines

The **ephemeral-port-range** command sets the starting port for the ephemeral port range. DataPower appliances use ephemeral ports to send data over TCP and UDP. To avoid conflicts between the ephemeral ports and the ports on which services listen, specify the starting port of the ephemeral port range. The last port in the ephemeral port range is 61000.

icmp-disable

This command disables the generation of a specific ICMP reply message.

Syntax

Disable ICMP reply messages

icmp-disable { **addressmask-reply** | **echo-reply** | **info-reply** | **timestamp-reply** }

Enables ICMP, if disabled

no icmp-disable { **addressmask-reply** | **echo-reply** | **info-reply** | **timestamp-reply** }

Parameters

addressmask-reply
echo-reply
info-reply
timestamp-reply

Specifies the target ICMP reply type.

Guidelines

The **icmp-disable** command disables the generation of a specific Internet Control Message Protocol (ICMP) reply message. By default, the appliance replies to the corresponding ICMP requests. Issue this command for each ICMP reply to disable.

Use the **no icmp-disable** command to enable the generation of a specific ICMP reply, if the ICMP reply is disabled. Issue this command for each ICMP reply to enable.

Examples

- Disable reply messages for ICMP echo (**ping**) requests.

  ```
  # icmp-disable echo-reply
  ```

- Enable **ping** reply messages to restore the default state.

  ```
  # no icmp-disable echo-reply
  ```

relax-interface-isolation

This command controls whether to allow packets from a wrong interface when both interfaces are on the same subnet.

Syntax

relax-interface-isolation { **on** | **off** }

Parameters

on Accepts a packet on an interface other than the one bound to the destination address of the packet. This setting is the default value.

off Allows only the interface that is bound to the destination address to accept the packet.

Guidelines

The **relax-interface-isolation** command controls whether to allow packets from a wrong interface when both interfaces are on the same subnet. As a security policy, the interface that receives a packet must be configured with the IP address that is the destination address of the packet. Enabling this option relaxes that restriction. The packet is allowed if the interface it arrives on contains an IP address in the same subnet as the destination address of the packet. Relax interface isolation, if destination-routing is enabled.

Example

Allows only the interface that is bound to the destination address to accept a packet.

```
# relax-interface-isolation off
```

reverse-path-filtering

This command determines whether incoming packets with a source address that cannot be routed by that interface are accepted and processed.

Syntax

reverse-path-filtering { **on** | **off** }

Parameters

on Ignores incoming packets with a source address that cannot be routed by that interface.

off Accepts and processes incoming packets with a source address that cannot be routed by that interface. This setting is the default value.

Guidelines

The **reverse-path-filtering** command determines whether incoming packets with a source address that cannot be routed by that interface are accepted and processed. Enabling this option effectively disables source routing because the appliance ignores such packets.

Note: If you allow reverse path filtering, the appliance cannot correctly route requests from the Sysplex Distributor to use the Target Control Server.

tcp-retries

This command sets the number of times to send a failed TCP SYN request.

Syntax

tcp-retries *retries*

Parameters

retries Specifies the number of times the local system sends a TCP SYN that receives no response. Enter a value in the range 1 - 32. The default value is 5.

Guidelines

The **tcp-retries** command sets the number of times the local system sends a failed TCP SYN request.

Example

Set to 10 attempts.

```
# tcp-retries 10
```

tcp-window-scale

This command determines whether to enable TCP window scaling.

Syntax

tcp-window-scale { **on** | **off** }

Parameters

on Enables TCP window scaling. This setting is the default value.

off Disables TCP window scaling.

Guidelines

The **tcp-window-scale** command determines whether to enable TCP window scaling. Window scaling allows the negotiation of window sizes greater than 64 KB. Disabling this option might help workaround TCP systems that do not understand or that misinterpret window scaling.

SSH client commands

You can use the SSH client commands to configure the client authentication on the IBM MQ Appliance.

The SSH client commands can be run from the command line interface in SSH client configuration mode. To enter SSH client configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:
   `config`
2. From global configuration mode, enter SSH client configuration mode:
   `sshclientprofile` *name*

   where *name* is the name of the configuration.
3. Type `exit` to leave the configuration mode and save your changes, then type `exit` again to leave global configuration mode.

client-known-host

This command adds or removes an SSH peer as an SSH known host.

Syntax

client-known-host *host* **ssh-rsa** *key*

no client-known-host *host*

Parameters

host Specifies the fully qualified host name or IP address for the peer. For example:

- `ragnarok.datapower.com`
- `10.97.111.108`

ssh-rsa
Identifies RSA as the key type.

key Specifies the host public key for the peer. For example:

```
AAAAB3NzaC1yc2EAAAABIwAAAIEA1J/99rRvdZmVvkaKvcG2a+PeCm25
p8OJl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfRl
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi0lud5Pl7+BXQ
VpPbyVujoHINCrx0k/z7Qpkozb4qZd8==
```

Guidelines

The **client-known-host** command adds an SSH peer as an SSH known host.

The **no client-known-host** command removes an SSH peer as an SSH known host.

Examples

- Add `ragnarok.datapower.com` by host name as an SSH client known host.

```
# client-known-host ragnarok.datapower.com ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEA1J/99rRvdZmVvkaKvcG2a+PeCm25
p8OJl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfRl
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi0lud5Pl7+BXQ
VpPbyVujoHINCrx0k/z7Qpkozb4qZd8==
#
```

- Add `ragnarok.datapower.com` by IP address as an SSH client known host.

```
# client-known-host 10.97.111.108 ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEA1J/99rRvdZmVvkaKvcG2a+PeCm25
p8OJl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfRl
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi0lud5Pl7+BXQ
VpPbyVujoHINCrx0k/z7Qpkozb4qZd8==
#
```

- Remove ragnarok.datapower.com by IP address as an SSH client known host.

```
# no client-known-host 10.97.111.108
#
```

persistent-connection-timeout

This command specifies the timeout value that a persistent connection can remain idle.

Syntax

persistent-connection-timeout *seconds*

Parameters

seconds

Specifies the timeout value. Enter a value in the range 1 - 86000. The default value is 120.

Guidelines

The **persistent-connection-timeout** command specifies the number of seconds a persistent connection is allowed to remain in an idle state before a timeout occurs that closes the persistent connection.

This command is valid when the SSH client profile connects to an SSH server.

persistent-connections

This command controls the negotiation of persistent connections for the SSH client profile.

Syntax

persistent-connections { **on** | **off** }

Parameters

on Enables the establishment of persistent connections. This setting is the default value.

off Disables the establishment of persistent connections.

Guidelines

The **persistent-connections** command specifies whether to enable or disable persistent connections. When enabled, persistent connections allows new SSH requests to reuse the connection of a previous SSH session without having to re-authenticate.

This command is valid when the SSH client profile connects to an SSH server.

profile-usage

This command specifies the type of the SSH client profile connection.

Syntax

Deprecated - Connects to a Secure Cloud Connection.

```
profile-usage scc
```

Connects to an SSH server.

```
profile-usage sftp
```

Parameters

scc Deprecated - The SSH Client profile connects to a Secure Cloud Connection.

sftp The SSH Client profile connects to an SSH server. This setting is the default value.

Guidelines

The **profile-usage** command specifies the type of the SSH client profile connection.

- Deprecated - If the profile is for a Secure Cloud Connection, use **scc**. Use of this value is valid for only appliances with the Application Optimization feature.
- If the profile is for the SFTP protocol, use **sftp**.

strict-host-key-checking

This command specifies how host keys are checked during the connection and authentication phase.

Syntax

strict-host-key-checking { **on** | **off** }

Parameters

on Rejects incoming SSH host keys from remote servers not in the known host list.

off Accepts incoming SSH host keys from remote servers not in the known host list. This setting is the default value.

Guidelines

The **strict-host-key-checking** command specifies how host keys are checked during the connection and authentication phase. By default, strict host key checking is disabled. When disabled the SSH client verifies the incoming host key against the keys in the known hosts list. If the host key does not match an existing known host entry for the remote server, the connection is rejected. If the known host list does not contain a host key for the remote server, the SSH client automatically accepts the host and adds its host key to the known host list.

When strict host key checking is enabled, the SSH client connects only to known hosts with valid SSH host keys that are stored in the known hosts list. Host keys not in the known host list are rejected.

This command is valid when the SSH client profile connects to an SSH server.

user-auth

This command specifies which SSH user authentication methods to allow for client authentication in the SSH client profile.

Syntax

user-auth *type*

Parameters

type Identifies the authentication method. Separate multiple methods with the plus sign (+) character.

publickey
Allows SSH public key for user authentication.

password
Allows SSH password for user authentication.

Guidelines

The **user-auth** specifies which SSH user authentication methods to allow for client authentication. Authentication can be one or more methods. At least one method is required. By default, both public key and password are allowed. With both methods defined, public key authentication is attempted first.

This command is valid when the SSH client profile connects to an SSH server.

user-name

This command specifies the user name for authenticating to the remote SSH server.

Syntax

user-name *user*

Parameters

user Specifies the user name for authenticating to the remote SSH server.

Guidelines

The **user-name** command specifies the user name for authenticating to the remote SSH server. The specified user name must match an existing user ID on the SSH server. Use with the **user-password** or **user-private-key** command.

Examples

- Set the user name to `testUser` with the password `pword` as the credentials for basic authentication on the remote SSH server.

```
# sshclientprofile profile1
New SSH Client Profile configuration
# user-name testUser
# user-password pword
```

- Set the user name to `testUser` with the `myKey` key as the credentials for basic authentication on the remote SSH server.

```
# sshclientprofile profile1
New SSH Client Profile configuration
# user-name testUser
# user-private-key myKey
```

user-password

This command specifies the password that is associated with the user name used for the SSH client.

Syntax

user-password *password*

Parameters

password
Specifies the password that is associated with the user name.

Guidelines

The **user-password** command specifies the password to use with the value of the **user-name** command.

This command is valid when the SSH client profile connects to an SSH server.

Example

Set the user name to `testUser` with the password `pword` as the credentials for basic authentication on the remote SSH server.

```
# sshclientprofile profile1
New SSH Client Profile configuration
# user-name testUser
# user-password pword
```

user-private-key

This command specifies the private key to use for public key authentication for the SSH client.

Syntax

user-private-key *alias*

Parameters

alias Specifies the name of a key alias.

Guidelines

The **user-private-key** command specifies the private key to use for public key authentication that is used with the value for the **user-name** command. A key that is used as a user private key must not be password-protected.

System Settings commands

System Settings mode provides the commands to modify system information for the appliance.

To enter the configuration mode, use the Global **system** command.

While in this mode, use the commands in the following table to define system settings.

- To view the current configuration, use the **show** command.
- To restore default values, use the **reset** command.
- To exit this configuration mode without saving changes to the running configuration, use the **cancel** command.

- To exit this configuration mode and save changes to the running configuration, use the **exit** command.

Table 38. System settings commands

Command name	Description
"**admin-state**" on page 437	This command sets the administrative state for the configuration.
"**contact**"	This command identifies the person or function responsible for appliance maintenance.
"**custom-ui-file**"	This command specifies the location of the custom user interface file.
"**entitlement**" on page 612	This command manages intrusion detection of the appliance in System Settings mode.
"**locale**" on page 612	This command specifies the locale for the operating language of the appliance.
"**location**" on page 613	This command specifies the location of the appliance.
"**name**" on page 613	This command sets an identifier for the appliance.

contact

This command identifies the person or function responsible for appliance maintenance.

Syntax

contact *contact*

Parameters

contact
Identifies the person or function responsible for appliance maintenance.

Guidelines

The **contact** command identifies the person who is responsible for managing the appliance. This information identifies the person who is responsible for managing this appliance by name, telephone number, email address, or a combination of these items.

Example

Set a system contact.

```
# contact "Tector Gorch; 555.555.5555; tector.gorch@example.com"
```

custom-ui-file

This command specifies the location of the custom user interface file.

Syntax

custom-ui-file *URL*

no custom-ui-file

Parameters

URL Specifies the location of the file on the appliance.

Guidelines

The **custom-ui-file** command specifies the location of the custom user interface file. The file must be in the `local:` or `store:` directory on the appliance. The file cannot be on a mounted file system.

This XML file contains custom user interface messages to display in the GUI and from the CLI. This file also defines the custom prompt for the CLI. After you create the custom user interface file, use the **test schema** command to validate that the XML file is conformant with the `dp-user-interface.xsd` schema.

Use the **no custom-ui-file** command to remove the use of custom messages and the CLI prompt that are defined in the custom user interface file.

Example

Specify the `xyzbanner.xml` file in the `store:` directory as the custom user interface file.

```
# custom-ui-file store:///xyzbanner.xml
```

entitlement

This command specifies the original serial number.

Syntax

entitlement *original*

Parameters

original
Specifies the original serial number.

Guidelines

The **entitlement** command specifies the serial number of the original appliance after you receive a replacement appliance. Without the serial number of the original appliance, IBM cannot entitle the replacement appliance for future maintenance or warranty service.

locale

This command specifies the locale for the operating language of the appliance.

Syntax

locale *locale*

Parameters

locale Specifies the operating language.

de German.

en English.

es Spanish.

fr French.

it Italian.

ja Japanese.

ko Korean.

pt_BR Brazilian Portuguese.

ru Russian.

zh_CN Simplified Chinese.

zh_TW Chinese (Taiwan).

Guidelines

The **locale** command specifies the locale for the operating language of the appliance. This setting manages locale-specific conventions, such as date and time formats, and controls the language of log messages.

Before a locale is selected, the administrative state of the language must be enabled with the **language** command.

location

This command specifies the location of the appliance.

Syntax

location *location*

Parameters

location
Specifies the appliance location.

Guidelines

The **location** command identifies the location of the appliance.

name

This command sets an identifier for the appliance.

Syntax

name *identifier*

Parameters

identifier
Specifies the identifier. Use a string up to 127 characters in length.

Guidelines

The **name** command specifies the system identifier of the appliance. Define an identifier that consists of only ASCII letters and numbers, for this value can be used to identify the appliance to establish a remote connection. Some protocols do not support spaces. If the name contains spaces, enclose in double quotation marks.

You should not specify a name that consists of only numbers.

Renaming appliances that belong to a high availability or disaster recovery configuration will prevent either from operating. See “Changing appliance names in high availability configurations” on page 130 and “Changing appliance names in disaster recovery configurations” on page 146.

When the custom user interface file defines the command-line extension, this identifier is added before the prompt.

Specify `no name` to clear a previously set name.

Examples

- Set `Duluth` as the name of the appliance.

  ```
  # name Duluth
  ```

- Set `Tango Lake` as the name of the appliance.

  ```
  # name "Tango Lake"
  ```

- Clear a previously set name:

  ```
  # no name
  ```

VLAN Commands

You can use the VLAN commands to configure the VLAN interfaces on the IBM MQ Appliance.

The VLAN commands can be run from the command line interface in VLAN configuration mode. To enter VLAN configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:

   ```
   config
   ```

2. From global configuration mode, enter Ethernet configuration mode:

   ```
   vlan name
   ```

   where *name* is the name of VLAN interface that you want to configure.

3. When you have finished, save your configuration:

   ```
   write memory
   ```

4. Type `exit` to leave the configuration mode, then type `exit` again to leave global configuration mode.

VLAN interfaces are not supported for links used in high availability configurations or disaster recovery configurations.

ethernet-interface

This command sets the Ethernet interface to provide connectivity.

Syntax

ethernet-interface {**eth10** | **eth11** | **eth12** | **eth13** | **eth14** | **eth15** | **eth16** | **eth17** | **eth21** | **eth22** | **mgt0** | **mgt1** }

Guidelines

The **ethernet-interface** command sets the Ethernet interface that provides connectivity to the VLAN interface. Even if the Ethernet interface is not configured with an IP address, this command enables that Ethernet port.

This command is meaningful only when you use the **over** command to indicate that connectivity is over an Ethernet interface.

Indicate that the `eth10` Ethernet interface is the parent.

```
# over ethernet
# ethernet-interface eth10
```

identifier

This command sets the VLAN identifier.

Syntax

identifier *number*

Parameters

number

Specifies the number of the VLAN identifier. Enter a value in the range 1 - 4094. The default value is 2.

Guidelines

The **identifier** command specifies the VLAN identifier to send traffic as well as to receive traffic. The identifier must be unique among all VLAN interfaces on the same Ethernet interface.

Set the VLAN identifier to 42.

```
# identifier 42
```

ip-address

This command assigns the primary network address for the VLAN interface.

Syntax

ip-address *address*

Parameters

address

Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

Guidelines

The **ip-address** command assigns the primary network address to the interface. The network address is an IP address with its subnet mask.

To assign secondary, or auxiliary, IP addresses, use the **ip-secondary-address** command.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

Examples

- Assign an IP address in version 4 format.

  ```
  # ip-address 192.168.7.6/27
  ```

- Assign an IP address in version 6 format.

  ```
  # ip-address 2001:0db8:3c4d:0015::abcd:ef12/34
  ```

ip-config-mode

This command identifies the configuration mode for the VLAN interface.

Syntax

ip-config-mode { **static** | **dhcp** | **slaac** }

Parameters

static Indicates a static, manual configuration. This setting is the default value.

dhcp Indicates IPv4 autoconfiguration with DHCP.

slaac Indicates IPv6 autoconfiguration with SLAAC.

Guidelines

The **ip-config-mode** command identifies the configuration mode of the interface.

- With the `static` keyword, define the configuration for the interface as provided by your network administrator.
  - Use the **ip-address** command to assign the primary network address.
  - Use the **ip-secondary-address** command to manage secondary, or auxiliary, network addresses.
  - Use the **ipv4-default-gateway** command to assign the default IPv4 gateway.
  - Use the **ipv6-default-gateway** command to assign the default IPv6 gateway.
  - Use the **ip-route** command to manage static routes in the routing table.
- With the `dhcp` keyword, the appliance ignores configuration data about the physical interface.
- With the `slaac` keyword, the appliance ignores configuration data about the physical interface.

Examples

- Change the configuration mode to IPv4 autoconfiguration with DHCP.

  ```
  # ip-config-mode dhcp
  ```

- Change the configuration mode to manual configuration.

  ```
  # ip-config-mode static
  ```

ip-route

This command manages static routes in the routing table for the VLAN interface.

Syntax

Add a static route

ip-route *address next-hop-address* [*metric*]

Delete a static route

no ip-route *address next-hop-address*

Parameters

address
Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

next-hop-address
Specifies the IP address of the next-hop router.

metric Optionally specifies the preference for the route. The lesser the value, the more preferred the route. For each IP family, the supported range differs.

- For IPv4, enter a value in the range 0 - 255. The default value is 0.
- For IPv6, enter a value in the range 0 - 65536. The default value is 512.

Guidelines

The **ip-route** command manages static routes in the routing table. Issue this command for each static route to add to the routing table.

To delete a static route, use the **no ip-route** command. Issue this command for each static route to delete from the routing table.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

Examples

- Add a static route to the routing table (subnet 10.10.10.224 via next-hop router 192.168.1.100). The metric for the route is 0, the default value for IPv4, which is the most preferred route.

  ```
  # ip-route 10.10.10.0/27 192.168.1.100
  ```

- Delete a static route from the routing table (subnet 10.10.10.224 via next-hop router 192.168.1.100).

  ```
  # no ip-route 10.10.10.0/27 192.168.1.100
  ```

ip-secondary-address

This command manages secondary network addresses for the VLAN interface.

Syntax

Add a secondary address
ip-secondary-address *address*

Remove a secondary address
no ip-secondary-address *address*

Remove all secondary addresses
no ip-secondary-address

Parameters

address
Specifies the IP address and netmask. The netmask is in CIDR format and is the integer that assigns the prefix length.

- For version 4, the prefix length can be in the range of 0 through 32.
- For version 6, the prefix length can be in the range of 0 through 128.

Guidelines

The **ip-secondary-address** command manages secondary network addresses for the current interface. The network address is the IP address and its subnet mask. A secondary IP address is a bind address. The secondary IP address is used only as a source IP address when it responds to incoming traffic to the secondary IP address.

To create the primary IP address, use the **ip-address** command.

To remove secondary IP addresses, use the **no ip-secondary-address** command.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

Examples

- Add 192.168.7.6/27 as a secondary IP address to the interface.

  ```
  # ip-secondary-address 192.168.7.6/27
  ```

- Remove 192.168.7.6/27 as a secondary IP address.

  ```
  # no ip-secondary-address 192.168.7.6/27
  ```

- Remove all secondary IP addresses.

  ```
  # no ip-secondary-address
  ```

ipv4-default-gateway

This command designates the default IPv4 gateway for the VLAN interface.

Syntax

Designates the default IPv4 gateway

ipv4-default-gateway *address*

Deletes the default IPv4 gateway

no ipv4-default-gateway

Parameters

address

Specifies the IP address of the default IPv4 gateway.

Guidelines

The **ipv4-default-gateway** command designates the default IPv4 gateway that the interface can reach. If the interface supports both IP families, use the **ipv6-default-gateway** command to designate the default IPv6 gateway.

Use the **no ipv4-default-gateway** command to delete the default IPv4 gateway.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

ipv6-dadtransmits

This command sets the number of IPv6 duplication address detection attempts for the VLAN interface.

Syntax

ipv6-dadtransmits *attempts*

Parameters

attempts
Specifies the number of attempts. The default value is 1.

Guidelines

The **ipv6-dadtransmits** command sets the number of IPv6 duplication address detection (DAD) attempts. This command is relevant for only IPv6 addresses on the appliance.

If you specify more than one attempt, use the **ipv6-nd-retransmit-timer** command to set the interval between attempts.

ipv6-default-gateway

This command designates the default IPv6 gateway for the VLAN interface.

Syntax

Designate the default IPv6 gateway
ipv6-default-gateway *address*

Delete the default IPv6 gateway
no ipv6-default-gateway

Parameters

address
Specifies the IP address of the default IPv6 gateway.

Guidelines

The **ipv6-default-gateway** command designates the default IPv6 gateway that the interface can reach. Define a default IPv6 gateway if you defined IPv6 IP addresses.

If the interface supports both IP families, use the **ipv4-default-gateway** command to designate the default IPv4 gateway.

Use the **no ipv6-default-gateway** command to delete the default IPv6 gateway.

This command is meaningful except when you use the **ip-config-mode** command for autoconfiguration with DHCP or SLAAC.

ipv6-nd-retransmit-timer

This command sets the interval between IPv6 neighbor discovery attempts for the VLAN interface.

Syntax

ipv6-nd-retransmit-timer *milliseconds*

Parameters

milliseconds
Specifies the interval between attempts in milliseconds. The default value is 1000.

Guidelines

The **ipv6-nd-retransmit-timer** command sets the interval neighbor discovery attempts. This command is relevant for only when the interface uses IPv6 addresses.

link-aggregation-interface

This command sets the aggregate interface to provide connectivity.

Syntax

link-aggregation-interface *name*

Parameters

name Specifies the name of an aggregate interface

Guidelines

The **link-aggregation-interface** command sets the aggregate interface to provide connectivity.

This command is meaningful only when you use the **over** command to indicate that connectivity is over an aggregate interface.

Indicate that the aggregate-1 aggregate interface is the parent.

```
# over link-aggregation
# link-aggregation-interface aggregate-1
```

mtu

This command sets the maximum transmission unit of the VLAN interface.

Syntax

mtu *bytes*

Parameters

bytes Specifies the maximum size in bytes. Enter a value in the range 576 - 16128. The default value is 1500.

Guidelines

The **mtu** command sets the maximum transmission unit (MTU) for the VLAN interface.

The MTU for the VLAN interface cannot be greater than the MTU of the parent interface. The parent interface is either an Ethernet interface or an aggregate interface. Use the **over** command to set the parent interface type.

- When the parent is an Ethernet interface, the **ethernet-interface** command sets the Ethernet interface. The MTU of the VLAN interface cannot be greater than the MTU of the Ethernet interface.
- When parent is an aggregate interface, the **link-aggregation-interface** command sets the aggregate interface. The MTU of the VLAN interface cannot be greater than the MTU of the aggregate interface.

The configuration for each interface type provides the **mtu** command. The MTU is determined regardless of the length of the layer 2 encapsulation.

Example

Set the MTU to 4 KB.

```
# mtu 4096
```

outbound-priority

This command sets the priority value in outbound packets.

Syntax

outbound-priority *priority*

Parameters

priority
Specifies the priority value. Enter a value in the range 0 - 7. The default value is 0.

Guidelines

The **outbound-priority** command sets the priority value in outgoing VLAN headers for packets. These packets are sent on this VLAN interface. This value is placed in the `user_control` field of the Tag Control Information (TCI). The exact interpretation of the value depends on the VLAN switch configuration.

Set the priority value to 4.

```
# outbound-priority 4
```

over

This command sets the parent interface type.

Syntax

over { **ethernet** | **link-aggregation** }

Parameters

ethernet
Indicate that the parent is an Ethernet interface. This setting is the default value.

link-aggregation
Indicates that the parent is an aggregate interface.

Guidelines

The **over** command sets the parent interface type for the VLAN interface.

- When an Ethernet interface, use the **ethernet-interface** command to set the parent Ethernet interface.
- When an aggregate interface, use the **link-aggregation-interface** command to set the parent aggregate interface.

Examples

Indicate that the eth10 Ethernet interface is the parent.

```
# over ethernet
# ethernet-interface eth10
```

Indicate that the aggregate-1 aggregate interface is the parent.

```
# over link-aggregation
# link-aggregation-interface aggregate-1
```

packet-capture

This command manages a packet-capture for the VLAN interface session.

Syntax

Start a packet-capture session

packet-capture *file seconds KB* ["*expression*"]

Stop a packet-capture session

no packet-capture *file*

Parameters

file Specifies the file name for the packet capture. You can simultaneously capture packets on multiple interfaces by specifying a different file name for each interface.

seconds

Specifies the maximum duration of the packet-capture session in seconds. Enter a value in the range 5 - 86400. The special value of -1 indicates that the packet capture is continuous and completes when it reaches the maximum file size or until you issue the **no packet-capture** command.

KB Specifies the maximum size of the file in KB. Enter a value in the range 10 - 500000.

expression

Optionally specifies the expression that filters the packet capture. Enclose the expression in double quotation marks.

Guidelines

The **packet-capture** command manages a packet-capture session on the current interface. The data from the session is saved in the pcap format. To interpret the packet, use a network protocol analyzer.

Use the **no packet-capture** command to stop a packet-capture session.

Examples

- Start a timed packet-capture session that writes data to the temporary:///capture-1 file. The session completes either after 30 minutes or when the file contains 2500 KB, whichever occurs first.

  ```
  # packet-capture temporary:///capture-1 1800 2500
  Trace begun.
  #
  ```

- Start a timed packet-capture session that writes data to the temporary:///capture-2 file. The session records only packets where 53 is the destination port. The session completes either after 30 minutes or when the file contains 2500 KB, whichever occurs first.

```
# packet-capture temporary:///capture-2 1800 2500 "dst port 53"
Trace begun.
#
```

- Start a continuous packet-capture session that writes data to the `temporary:///capture-3` file. The session completes either when it contains 50000 KB or when you stop it.

```
# packet-capture temporary:///capture-3 -1 50000
Trace begun.
#
```

- Stop the packet-capture session that writes data to the `temporary:///capture-3` file.

```
# packet-capture temporary:///capture-3
Continuous packet capture to temporary:///capture-3 on interface stopped.
#
```

Web management service commands

You can use the web management service commands to configure the web management service settings on the IBM MQ Appliance.

The web management service commands can be run from the command line interface in web management service configuration mode. To enter web management service configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:

   ```
   config
   ```

2. From global configuration mode, enter throttle configuration mode:

   ```
   web-mgmt
   ```

3. Type `exit` to leave the configuration mode and save your changes, then type `exit` again to leave global configuration mode.

While in this mode, use the commands in the following table to modify the web management service.

idle-timeout

This command specifies the idle-session timer for the web management service.

Syntax

idle-timeout *seconds*

Parameters

seconds

Specifies the timeout value of the idle session in seconds. Use any value of 0 - 65535. The default value is 600. A value of 0 disables the timer.

local-address

This command changes the local address and port that the web management service monitors for requests.

Syntax

local-address *address port*

Parameters

address
Specifies the IP address or host alias of a local interface that the service listens on. The default value is `0.0.0.0`.

port Specifies the listening port for the service. The default value is 9090.

Guidelines

The **local-address** command specifies the local address and port that the service listens on.

- When the value is `0.0.0.0`, the service listens on all active IPv4 addresses.
- When the value is `::`, the service listens on all active IPv4 and IPv6 addresses.

Attention: If the service is a management service, the value of `0.0.0.0` or `::` is a security risk.

If the local address supports IPv6 addresses, modify the `web-mgmt` ACL to include an `allow` clauses for specific or all IPv6 addresses.

Use a host alias to help ease migration tasks among appliances. To create a local host alias, use the Global **host-alias** command.

To change the port, use the **port** command.

Example

Specify that port 8090 on all interfaces is monitored for requests to the Web Management Service.

```
# web-mgmt
Web Management Service configuration mode
# local-address 0.0.0.0 8090
```

port

This command changes the listening port for the web management service.

Syntax

port *port*

Parameters

port Indicates the listening port. The default value is 9090.

Examples

Change the port to 8090 on the local interfaces that monitor for requests to the web management service.

```
# web-mgmt
Web Management Service configuration mode
# port 8090
#
```

ssl (deprecated)

This command changes the association of the SSL proxy profile for the web management interface.

Syntax

ssl *name*

Parameters

name Specifies the name of an SSL proxy profile.

Guidelines

The **ssl** command changes the association of the SSL proxy profile for the web management interface. You must associate an SSL proxy profile with the web management interface. By default, the web management interface uses the shipped configuration. Unless you create and assign a different SSL proxy profile, all appliances use the same SSL configuration.

To create an SSL proxy profile, use the Global **sslproxy** command.

ssl-config-type

This command sets the type of the SSL profile for the web management interface.

Syntax

ssl-config-type {proxy | server | sni}

Parameters

proxy (deprecated)

Uses an SSL proxy profile with the cryptographic profiles to secure connections. This setting is the default value for backward compatibility.

server

Uses an SSL server profile to secure connections.

sni

Uses an SSL SNI server profile to secure connections.

Guidelines

The **ssl-config-type** command sets the SSL profile type to secure connections between clients and the appliance. You can use an SSL proxy profile, an SSL server profile, or an SSL SNI server profile.

- The SSL proxy profile is deprecated. Check whether your configuration uses an SSL proxy profile. If yes, modify your configuration to use an SSL server profile or an SSL SNI server profile to secure connections.
- When the appliance is an SSL server, you use an SSL server profile to secure connections. To specify an SSL server profile, use the ssl-server command.
- When the appliance is an SSL server and supports Server Name Indication (SNI), you use an SSL SNI server profile. To specify an SSL SNI server profile, use the ssl-sni-server command.

ssl-server

This command associates an SSL server profile with the web management interface.

Syntax

ssl-servername *name*

Parameters

name

Specifies the name of an SSL server profile.

Guidelines

The **ssl-server** command specifies the SSL server profile to secure connections between clients and the appliance. You use an SSL server profile when the appliance is an SSL server.

To create an SSL server profile, use the Crypto **ssl-server** command.

This command is relevant when the type set by the **ssl-config-type** command is `server`.

user-agent

This command associates a user agent with the web management service.

Syntax

user-agent *name*

no user-agent

Parameters

name Specifies the name of a user agent.

Guidelines

The **user-agent** command associates a user agent. The user agent must exist in the current domain. You can associate only one user-agent.

To create a user agent, use the Global **user-agent** command.

To remove the association of the user agent, use the **no user-agent** command.

Messages

Visit the messages section in the IBM MQ documentation for help with IBM MQ diagnostic messages.

http://www.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.ref.doc/q050250_.htm

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.

© Copyright IBM Corp. 2015, 2016

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.