

Version 9 Release 1

IBM MQ Appliance

IBM

Note

Before using this information and the product it supports, read the information in [“Notices” on page 1209](#).

This edition applies to version 9 release 1 modification 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 International Business Machines Corporation 2015, 2020.**

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

Contents

- Product overview..... 11**
 - Introduction to the IBM MQ Appliance..... 11
 - Appliances and the appliance firmware.....11
 - Relationship with IBM DataPower appliances..... 12
 - Appliance comparison: M2002 and M2001.....12
 - IBM MQ Appliance product identifiers and export information..... 13
 - What's new and changed in release 9.1.0..... 13
 - What's new and changed in release 9.1.1..... 14
 - What's new and changed in release 9.1.2..... 14
 - What's new and changed in release 9.1.3..... 15
 - What's new and changed in release 9.1.4..... 15
 - What's new and changed in release 9.1.5..... 15
 - High availability.....16
 - Disaster recovery for a high availability configuration..... 18
 - Disaster recovery..... 19
 - Accessibility features for IBM MQ Appliance20
 - Notices..... 20
 - Trademarks.....22
 - Terms and conditions for product documentation.....23
 - IBM Online Privacy Statement.....23
- Planning.....25**
 - Differences between administering an IBM MQ Appliance and an IBM MQ installation..... 25
 - Control commands on the IBM MQ Appliance..... 25
 - Differences between queue managers that are running on the IBM MQ Appliance and an IBM MQ installation..... 31
 - Exits and services on the IBM MQ Appliance..... 31
 - Queue manager configuration on the IBM MQ Appliance.....32
 - IBM MQ objects on the IBM MQ Appliance..... 33
 - XA transactions..... 33
 - Planning a high availability system..... 33
 - Network requirements for high availability..... 35
 - Planning a disaster recovery system.....36
 - Hardware limitations for disaster recovery..... 37
 - Planning network connections..... 38
 - Network connections on an M2002 appliance.....38
 - Network connections on an M2001 appliance.....42
 - Network configuration guidance.....46
 - Planning SAN storage..... 51
 - Capacity planning.....52
 - Security planning..... 53
 - IBM MQ and IBM MQ Appliance on premises considerations for GDPR readiness..... 54
- Installing the M2001 appliance..... 65**
 - Safety..... 65
 - Guidelines for servicing electrical equipment..... 66
 - Inspecting for unsafe conditions..... 67
 - Safety statements..... 67
 - Introducing the IBM MQ Appliance M2001..... 70
 - Specifications and features.....70
 - Intrusion detection.....72
 - Components on the front.....72
 - Components on the rear..... 79

Prepare for installation.....	80
Rack requirements.....	80
Tool requirements.....	81
Installing the appliance in a rack.....	82
Installing rails in the rack frame.....	82
Installing the appliance on the rails.....	85
Connect the appliance to an AC power source.....	86
Connect the appliance to your network.....	87
Setting up the initial firmware configuration.....	88
Configuration requirements.....	88
Firmware considerations.....	89
Procedure 1 of 3: Connecting the serial cable to the appliance.....	90
Procedure 2 of 3: Initializing the appliance.....	90
Procedure 3 of 3: Accepting the license agreement.....	92
Maintenance.....	92
Diagnosing your appliance.....	92
Troubleshooting your appliance.....	99
Removing or replacing the appliance or parts.....	101
Removing the batteries.....	117
Installing the M2002 appliance.....	119
Safety.....	119
Guidelines for servicing electrical equipment.....	120
Inspecting for unsafe conditions.....	121
Safety statements.....	121
Labels.....	124
Introducing the IBM MQ Appliance M2002.....	124
Specifications and features.....	124
Intrusion detection.....	126
Identifying components.....	127
Installation preparation.....	134
Installation of the appliance in a rack.....	134
Installing rails in the rack frame.....	135
Installing the appliance on the rails.....	138
Considerations when connecting the appliance to AC power sources.....	140
Considerations when connecting the appliance to your network.....	140
Setting up the initial firmware configuration.....	142
Configuration requirements.....	142
Installation considerations.....	142
1 of 3: Connecting the serial cable to the appliance.....	143
2 of 3: Initializing the appliance.....	144
3 of 3: Accepting the license agreement.....	145
Completing the configuration.....	146
Diagnostics and the appliance.....	146
M2002 LED diagnostics.....	146
Hardware testing from the CLI.....	148
Running the diagnostic self-test.....	149
Status providers for sensors.....	150
Troubleshooting your appliance.....	151
Troubleshooting fan modules.....	152
Troubleshooting power supply modules.....	152
Troubleshooting solid-state disk drive modules.....	153
Troubleshooting the appliance.....	153
Removing or replacing the appliance or parts.....	153
Installation guidelines.....	153
Part listings.....	154
Turning off the appliance.....	157
Replacing a fan module.....	158

Replacing a power supply module.....	159
Replacing a solid-state disk drive module.....	161
Replacing an Ethernet module.....	163
Replacing an SFP+ or a QSFP+ transceiver.....	166
Removing the appliance from the rack.....	168
Replacing an appliance.....	169
Removing batteries.....	170
Returning an appliance or part.....	172
Upgrading and downgrading firmware.....	173
Installing new firmware.....	173
Installing a new level of firmware by using the command line.....	174
Installing a new level of firmware by using the IBM MQ Appliance web UI.....	175
Reverting firmware.....	175
Reverting to the previous level of firmware by using the command line.....	176
Reverting to the previous level of firmware by using the IBM MQ Appliance web UI.....	176
Downgrading.....	176
Suspending an appliance from an HA group for maintenance.....	177
Upgrading a version 8.0 IBM MQ Appliance to version 9.1.....	177
Upgrading to version 9.1 by using the command line.....	178
Upgrading to version 9.1 by using the IBM MQ Console.....	179
Upgrading a Version 8.0 HA configuration to version 9.1.....	180
Upgrading a Version 8.0 HA configuration with DR to version 9.1.....	181
Upgrading a Version 8.0 DR configuration to version 9.1.....	182
Upgrading a version 9.0 IBM MQ Appliance to version 9.1.....	183
Upgrading to version 9.1 by using the command line.....	184
Upgrading to version 9.1 by using the IBM MQ Console.....	185
Upgrading a Version 9.0 HA configuration to version 9.1.....	186
Upgrading a Version 9.0 HA configuration with DR to version 9.1.....	187
Upgrading a Version 9.0 DR configuration to version 9.1.....	188
Transferring queue managers to other IBM MQ Appliances.....	189
Transfer from an existing single appliance to a new single appliance by using archive files.....	189
Transfer from an existing single appliance to a new single appliance by using DR commands.....	189
Transfer from an existing appliance in a disaster recovery configuration.....	193
Transfer from an existing high availability pair of appliances to a new pair of appliances.....	198
Configuring.....	207
Command line access.....	207
Connecting to the serial port.....	207
Configuring the SSH service.....	207
Configuring the IBM MQ Appliance web UI.....	209
Changing the IBM MQ Appliance web UI IP address and port.....	210
Configuring certificates for IBM MQ Appliance web UI.....	211
Customizing the user interfaces.....	214
Supported markup for the user interface customization file.....	214
Template of the custom user interface file.....	216
Configuring the appliance.....	217
Ethernet interfaces.....	217
VLAN interfaces.....	220
Link aggregation interfaces.....	222
IPMI Settings.....	224
DNS settings.....	225
SNMP Settings.....	227
Configuring the locale, date, and time.....	238
Configuring the appliance name.....	242
Configuring appliance user access.....	243
Configuring the REST management interface.....	244
Configuring the appliance by using the REST management interface.....	248
Retrieving configuration information by using REST.....	248

Modifying and deleting existing configurations by using REST.....	250
Creating configurations by using REST.....	252
Configuring user access to the IBM MQ Console and the CLI.....	254
Granting full administrative access to appliance system settings and IBM MQ.....	255
Granting full administrative access to appliance system settings but barring access to IBM MQ.....	256
Granting access to IBM MQ but barring access to the appliance system settings	257
Granting limited access to a queue manager.....	259
Configuring queue managers.....	260
Configuring uniform clusters - appliance-specific considerations.....	261
Configuring environment variables.....	261
Adding an environment variable.....	262
Modifying an environment variable.....	263
Removing an environment variable.....	264
Viewing environment variables.....	264
Configuring IBM MQ Advanced Message Security.....	265
Configuring MCA interception.....	266
Configuring high availability.....	267
Configuring the hardware for high availability.....	267
Configuring the high availability group.....	271
Configuring high availability queue managers.....	275
Configuring error reporting for high availability.....	281
Example network set up for HA configuration.....	281
Configuring disaster recovery for a high availability queue manager.....	282
Replacing a failed node in a high availability group.....	286
Configuring disaster recovery.....	288
Configuring the hardware for disaster recovery.....	288
Configuring disaster recovery queue managers.....	291
Configuring disaster recovery for a high availability queue manager.....	296
Configuring SAN storage.....	300
Configuring SAN for the appliance.....	301
Configuring volumes.....	302
Initializing the file system for a volume.....	303
Configuring queue managers to use SAN storage.....	304
Removing queue managers that use SAN storage.....	305
Configuring the IBM MQ Console and REST API.....	306
Configuring logging for administrative REST API and IBM MQ Console.....	306
Configuring the LTPA token expiry interval.....	307
Configuring the response timeout.....	308
Configuring CORS for the REST API.....	308
Administering.....	311
Using the IBM MQ Console.....	311
Working with queue managers.....	311
Working with IBM MQ objects.....	315
Working with authority records.....	329
Monitoring system resource usage.....	332
Configuring dashboard layouts.....	344
The dashboard controls.....	346
Keyboard shortcuts.....	346
Administering IBM MQ by using the REST API.....	346
Message queue control commands.....	347
Administering messaging users.....	347
Using MQSC commands.....	349
Using an IBM MQ client	349
Setting up a queue manager to accept client connections.....	349
Configuring queue managers and objects by using a client.....	352
Putting and getting messages.....	352
Publishing and subscribing.....	353

Browsing a message queue.....	353
Creating and downloading a CCDT file.....	354
Starting and stopping the appliance.....	356
Restarting the appliance.....	356
Shutting down the appliance.....	357
Restarting queue managers by using the command line.....	357
Back up and restore.....	358
Backing up or saving the appliance configuration.....	358
Restoring the appliance configuration.....	359
Backing up messaging users.....	360
Restoring messaging users.....	361
Backing up a key repository.....	363
Restoring a key repository.....	364
Backing up a queue manager.....	365
Restoring a queue manager.....	367
Backing up IBM MQ Appliance web UI configuration data.....	367
Restoring IBM MQ Appliance web UI configuration data.....	368
Factory reset.....	369
Triggering appliance operations by using the REST management interface.....	370
Operating in a high availability environment.....	373
Suspending an appliance from an HA group for maintenance.....	373
Replacing a failed node in a high availability group.....	375
Managing queue manager locations in a high availability group.....	376
Viewing the status of appliances in a high availability group.....	377
Viewing the status of a high availability queue manager.....	378
Regenerating the keys for secure communication of the HA pair.....	380
Disaster recovery for a high availability queue manager.....	380
Operating in a disaster recovery environment.....	381
Switching over to a recovery appliance.....	381
Switching back to the main appliance.....	382
Replacing a failed node in a disaster recovery configuration.....	392
Replacing failed high availability nodes in a disaster recovery configuration	395
Testing the recovery appliance.....	399
Reversing disaster recovery roles.....	400
Viewing the status of a disaster recovery queue manager.....	401
Managing files by using the IBM MQ Appliance web UI.....	402
Managing files by using the REST management interface.....	405
Watchdog timer.....	409
Migrating and consolidating.....	411
Moving queue managers from other IBM MQ platforms.....	411
Moving a queue manager.....	411
Moving queue managers secured by using TLS.....	414
Planning for incompatible features in the queue manager.....	415
Handling incompatible features in the queue manager.....	416
Editing qm.ini files.....	417
Security.....	423
Types of user and how they are authenticated.....	423
User authorization, credential mapping, and access profiles.....	423
Access policies.....	425
Role based management.....	439
Important: avoiding user lock out when configuring role based management.....	440
User authentication with LDAP.....	440
User authentication with XML file.....	455
User authentication with local users.....	459
Credential mapping with an XML file.....	460
Credential mapping with local user groups.....	464
Password policy.....	467

Account policy.....	470
Local users and user groups.....	471
Routine user administration.....	476
Changing your own password by using the IBM MQ Appliance web UI.....	476
Changing your own password by using the command line.....	476
Resetting a user's password by using the IBM MQ Appliance web UI.....	477
Resetting a user's password by using the command line.....	477
Forcing a password change by using the IBM MQ Appliance web UI.....	478
Forcing a password change by using the command line.....	478
Resetting failed login count by using the IBM MQ Appliance web UI.....	478
Resetting failed login count by using the command line.....	479
TLS certificate management.....	479
Working with self-signed certificates.....	480
Working with CA-signed certificates.....	482
Listing certificates for a queue manager.....	487
Viewing a certificate for a queue manager.....	487
Deleting a certificate.....	488
Managing certificates on the appliance.....	488
FIPS compliance.....	494
Audit trail.....	494
Monitoring and reporting.....	497
IBM MQ appliance logs.....	497
IBM MQ error logs.....	499
IBM MQ diagnostic message services.....	505
Appliance audit logs.....	512
Appliance log targets.....	516
Monitoring system resource usage.....	522
Monitoring system resource usage by using the status command.....	522
Monitoring system resource usage by using the amqsrua command.....	523
Monitoring system resource usage by using the IBM MQ Console	526
Monitoring the appliance by using the show command.....	536
Developing your own resource monitoring program.....	538
Monitoring queue manager health.....	539
Monitoring the health of high availability and disaster recovery configurations.....	540
HA notification examples.....	540
DR notification examples.....	544
Application activity trace.....	547
Subscriptions to application activity trace.....	547
Creating subscriptions to application activity trace.....	547
Application activity trace: subscriptions compared with central collection.....	549
Using amqsact to view trace messages.....	549
Configuring trace levels.....	551
System topics for monitoring and activity trace.....	552
Monitoring the appliance by using the REST management interface.....	553
Example of retrieving status by using REST.....	553
Monitoring the IBM MQ Appliance by using SNMP.....	556
Troubleshooting.....	557
Appliance error reports.....	557
Configuring failure notification using the web UI.....	558
Generating error reports.....	559
Viewing error reports.....	559
Sending error reports.....	560
Using trace.....	560
Using trace in the IBM MQ Console.....	561
Resolving a partitioned problem in a high availability configuration.....	561
Resolving a partitioned problem in a disaster recovery configuration.....	564
Resolving an HA queue manager left in an indeterminate state	565

Troubleshooting file copy.....	566
Troubleshooting SAN problems.....	567
Problems resizing queue managers.....	569
Help with using runmqras	569
Recovering from hardware failures, M2001 model.....	569
Appliance fails, both disks unaffected.....	570
Appliance fails, one disk unaffected.....	571
Appliance operational, one disk in RAID pair fails.....	571
Recovering from hardware failures, M2002 model.....	572
Appliance fails, all disks unaffected.....	573
Appliance fails, three disks unaffected.....	574
Appliance operational, one disk in RAID set fails.....	575
Reference.....	577
Command reference.....	577
IBM MQ commands.....	577
Appliance commands.....	707
REST management interface.....	1000
REST request structure.....	1000
REST response structure.....	1001
REST management resources summary.....	1002
Query parameters.....	1005
REST API resources.....	1006
Messages.....	1208
Notices.....	1209
Programming interface information.....	1210
Trademarks.....	1210

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 V9.1 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 V9.1 network by using TLS.

You run applications on clients that connect to the 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. The appliance supports the same CCSIDs as Linux®.

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 217](#).

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 311](#). 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 [“Migrating and consolidating” on page 411](#).

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 16](#), and the disaster recovery feature, see [“Disaster recovery” on page 19](#).

Relationship with IBM DataPower appliances

The IBM MQ Appliance includes components from the IBM DataPower Gateway family of products.

The IBM MQ Appliance firmware builds on the long-term expertise that IBM has in developing network appliances by including components from the IBM DataPower Gateway family of products. This relationship is visible in a number of places, for example, you will see mentions of 'DPOS' (the low level firmware/operating system of the appliance) in the system logs. The IBM MQ Appliance is, however, a discrete and stand-alone product; there is much functionality from IBM DataPower Gateway appliances that is not present in the IBM MQ Appliance, and vice versa.

You do not need to be familiar with IBM DataPower to work with the IBM MQ Appliance. The information that you need is supplied in this documentation, and day to day management tasks are intuitive, whichever interface you choose to use. However, in many cases you will find that any experience with DataPower appliances (for example, first-time setup, working with the CLI, web UI, and other management interfaces) is valuable and carries over to the IBM MQ Appliance.

Appliance domains

One fundamental aspect of the IBM DataPower appliance that is not carried over to the IBM MQ Appliance is the concept of an 'Application Domain'. In DataPower products, domains provide a mechanism for the separation of applications from unrelated areas (Lines of Business, Test versus Production, and so on). In some senses, this is similar to the separation provided by connecting applications to a different queue manager hosted on the same appliance.

Therefore, the domain feature is not currently used in the IBM MQ Appliance, and creation of new domains is disabled. However, by the very nature of the platform, you will come across 'domains' mentioned in a few contexts, for example, in creating REST URIs, or displaying system objects. Therefore, for the purposes of DPOS as exploited in the IBM MQ Appliance, only the 'default' domain is required, and 'default' should always be supplied in these contexts.

Appliance comparison: M2002 and M2001

The following table 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](#)

M2002 (current model)	M2001 (previous model)
2 x 12 core "Skylake" Intel Xeon Gold 6126T processors (6 cores active in B model). 2.60 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, Twelve 2666 MHz DDR4 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
6 x 10 Gb Ethernet ports	4 x 10 Gb Ethernet ports
4 x 40 Gb Ethernet ports	-
4 x 3.2 TB SSDs under hardware controlled mirrored RAID (6 TB effective, 3 TB effective B model)	2 x 3.2 TB SSDs under hardware controlled mirrored RAID (3 TB effective)
2 GB RAID cache	1 GB RAID cache
SAN not supported	Fibre channel adapter for SAN connectivity

For more information on hardware specification, see [“Introducing the IBM MQ Appliance M2001”](#) on page 70 and [“Introducing the IBM MQ Appliance M2002”](#) on page 124.

IBM MQ Appliance product identifiers and export information

An overview of IBM MQ Appliance product offerings, their associated product identifiers (PIDs), and export classifications.

IBM MQ product offering	Product identifier (PID)	Machine-Type Model (MTM)	Firmware Export classification (ECCN)	Hardware Export classification (ECCN)
IBM MQ Appliance M2000	5725-S14	8436-54X	5D992	5A992
IBM MQ Appliance M2001	5725-Z09	8436-55X	5D992	5A992
IBM MQ Appliance M2002	5737-H47	8441-5	5D992	5A002.a.2

► V 9.1.0 What's new and changed in release 9.1.0

This topic describes new and changed features in version 9.1.0 of the appliance firmware.

The following features are new for version 9.1.0:

- A new appliance model, the M2002, is now available, see [“Appliance comparison: M2002 and M2001”](#) on page 12 and [“Introducing the IBM MQ Appliance M2002”](#) on page 124.
- For features that are new and changed with IBM MQ V9.1.0, see [What's new and changed in IBM MQ V9.1.0](#) in the IBM MQ documentation.

The following features are changed for Version 9.1.0:

- You can now specify custom interfaces for replication in high availability (HA) and disaster recovery (DR). This allows you, for example, to use aggregated links or the 40 Gb links on the M2002 appliance

for data replication. See [“Configuring custom HA replication interfaces”](#) on page 269 and [“Configuring custom DR replication interfaces”](#) on page 290.

- If you are upgrading a Version 8.0 or Version 9.1 high availability (HA) configuration, a manual step of upgrading the queue managers is required to configure the HA group in preparation for the change to Version 9.1. See [“Upgrading a Version 8.0 HA configuration to version 9.1”](#) on page 180 and [“Upgrading a Version 9.0 HA configuration to version 9.1”](#) on page 186 for details.

What's new and changed in release 9.1.1

This topic describes new and changed features in version 9.1.1 of the appliance firmware.

The following features are changed for Version 9.1.1:

- Block non-management traffic feature removed. By default, the appliance previously blocked non-management traffic on all network interfaces when it detected that one such interface was not correctly configured. It was recommended that you cleared that feature, particularly if you were configuring high availability. From Version 9.1.1 the block non-management traffic feature is removed so no action needs to be taken.
- The configuration scripts run by the **exec** command can now include **mqcli** commands. Note that the configuration file flow logic (implemented with the if/endif statements) is not available for **mqcli** commands.

What's new and changed in release 9.1.2

This topic describes new and changed features in version 9.1.2 of the appliance firmware.

The following features are new for Version 9.1.2:

- Queue manager error log enhancements, including support for:
 - Configuring the error log size
 - Excluding and suppressing messages
 - File-based diagnostic message services, including JSON formatted logs
 - Appliance log target integration, which provides support for streaming queue manager error log messages to a remote syslog server.

See [“IBM MQ appliance logs”](#) on page 497.

- The IBM MQ level name is now reported by the `show version` and `show firmware-version` appliance commands, see [“show version”](#) on page 911 and [“show firmware-version”](#) on page 892.
- The firmware version information and the CLI login message now identify whether the firmware delivery type is long-term support (LTS) or continuous delivery (CD).
- A controlled shutdown is now performed for queue managers during shutdown or reboot of the appliance.

The following features are changed for Version 9.1.2:

- The default values for some log target properties have changed. The default log target type has changed from `cache` to `file`, and the default timestamp format has changed from `syslog` to ISO Coordinated Universal Time (UTC) format (`zulu`).
- The maximum SSH banner length (before truncation) has increased from 1023 characters to 4096 characters.
- In the CLI, the IBM MQ administration mode (`mqcli`) is now also available within the global configuration command shell (`config`) in addition to the initial login shell. See [“IBM MQ commands”](#) on page 577.
- The REST URI for the `show firmware` command has changed from `/mgmt/status/default/FirmwareStatus` to `/mgmt/status/default/FirmwareStatus2`. The previous REST URI is still available, but is now deprecated. See [“/mgmt/status/default/FirmwareStatus2”](#) on page 1061.

- The REST URI for the **show firmware-version** command has changed from `/mgmt/status/default/FirmwareVersion2` to `/mgmt/status/default/FirmwareVersion3`. The previous REST URI is available, but is now deprecated. See [“/mgmt/status/default/FirmwareVersion3”](#) on page 1065.
- Public certificates are no longer provided in the `pubcert:` directory. You should request and upload any public certificates that you require. See [“Uploading certificates to the appliance”](#) on page 489.
- The appliance now runs IBM MQ V9.1.2. See [What's new and changed in IBM MQ Version 9.1.2](#) IBM MQ documentation for details of features that are new for Version 9.1.2.

What's new and changed in release 9.1.3

This topic describes new and changed features in version 9.1.3 of the appliance firmware.

The following features are new for Version 9.1.3:

- Status notifications are now issued when HA and DR queue managers, and the HA group, changes state. You can use these notifications to track the health of the HA or DR configuration. See [“Monitoring the health of high availability and disaster recovery configurations”](#) on page 540 for details.
- Support has been added for converting between CCSIDs 37 and 500 on the appliance.

The following features are changed for Version 9.1.3:

- The appliance now runs IBM MQ V9.1.3. See [What's new and changed in IBM MQ Version 9.1.3](#) IBM MQ documentation for details of features that are new for Version 9.1.3.

V 9.1.4 What's new and changed in release 9.1.4

This topic describes new and changed features in version 9.1.4 of the appliance firmware.

The following features are new for Version 9.1.4:

- You can now discover the operational mode of the appliance (for example, M2002A, M2002B, M2002B+) by using the **show system** command or by using a REST API command to show system settings. See [“Appliance modes”](#) on page 89 for M2001 appliances and [“Installation considerations”](#) on page 142 for M2002 appliances.
- IBM MQ events are now recorded in the appliance audit and CLI logs, see [“Appliance audit logs”](#) on page 512.
- A new parameter for the **test tcp-connection** command enables you to specify a local address to use when testing connectivity to a remote system. This enhancement enables network routing and firewall rules to be verified for a secondary IP address, such as a high availability floating IP address. See [“test tcp-connection”](#) on page 717.
- New parameters have been added to the **crtmqm**, **setmqini**, and **runmqsc** commands to support the automatic configuration of queue managers that are members of uniform clusters. See [“Configuring uniform clusters - appliance-specific considerations”](#) on page 261.

The following features are changed for Version 9.1.4:

- HA Queue managers in the primary role or secondary role are now reported as having the role none when an appliance is suspended. See [“HA notification examples”](#) on page 540.
- The appliance now runs IBM MQ V9.1.4. See [What's new and changed in IBM MQ Version 9.1.4](#) IBM MQ documentation for details of features that are new for Version 9.1.4.

V 9.1.5 What's new and changed in release 9.1.5

This topic describes new and changed features in version 9.1.5 of the appliance firmware.

The following features are new for Version 9.1.5:

- You can now write appliance error reports to the `mqdiag` directory on the RAID disk. See [“Appliance error reports”](#) on page 557.

- There is a new timeout feature to use when specifying an ntp service for the appliance to use. The timeout specifies the time that connection to the NTP server is attempted for. See [“Configuring NTP service settings by using the IBM MQ Appliance web UI”](#) on page 238 and [“Configuring NTP service settings by using the command line”](#) on page 238.
- A new command, **reset-ssh-keys**, can be used to delete old SSH keys, see [“reset-ssh-keys”](#) on page 812.
- The **dspmqttrc** command can now be run on the appliance to format trace files. See [“dspmqttrc \(display formatted trace\)”](#) on page 593.
- You can now control the size of queue files. You have the option of configuring and monitoring queues that can support substantially more than the two terabyte default limit used in releases of before IBM MQ Version 9.1.5. You also have the option of reducing the size a queue file can grow to. To enable you to configure queues, there is an additional attribute on local and model queues, MAXFSIZE, and to monitor queues there are two additional queue status attribute, CURFSIZE and CUR|AXFS. See [Modifying IBM MQ queue files in the main IBM MQ documentation](#).

The following features are changed for Version 9.1.5:

- You can now configure a disaster recovery solution for a high availability group without specifying a floating IP address. This removes the requirement for both HA appliances being in the same subnet for configuring disaster recovery. See [“Configuring disaster recovery for a high availability queue manager”](#) on page 282.

Note: You can currently configure only one HA/DR queue manager with no floating IP address, contact IBM Support for help with overcoming this limitation.

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 18. 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.

Applications can connect to HA queue managers in one of two ways. They can have an IP address configured for the data interface on each of the appliances in the HA group, and the application itself determines which one to use for connecting to the active queue manager. Alternatively, applications can use a single floating IP address to access a particular queue manager, and that IP address will work for that queue manager whichever appliance it is running on (note that the appliances still both require an interface configured with a static IP for the floating IP address to map to). Using a floating IP address in this way makes queue manager failover almost invisible to the connecting application.

For more information about configuring HA on the IBM MQ Appliance, see [“Configuring high availability”](#) on page 267. For information about planning the physical location, see [“Planning a high availability system”](#) on page 33.

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.

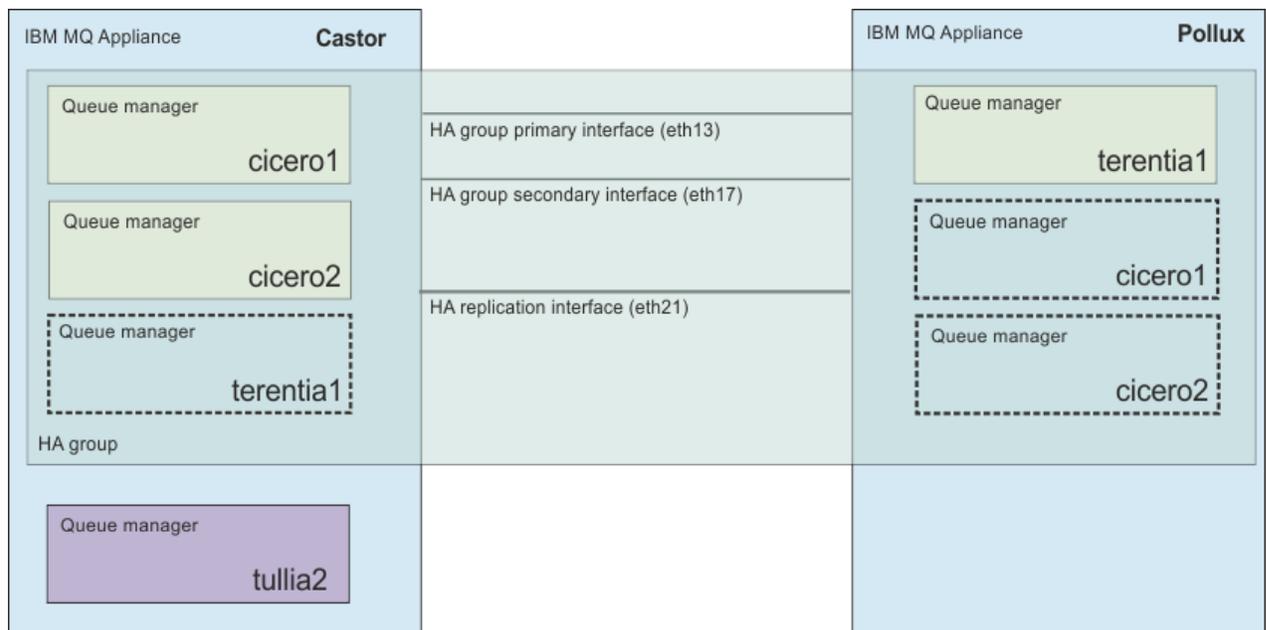


Figure 1. Example HA group

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.

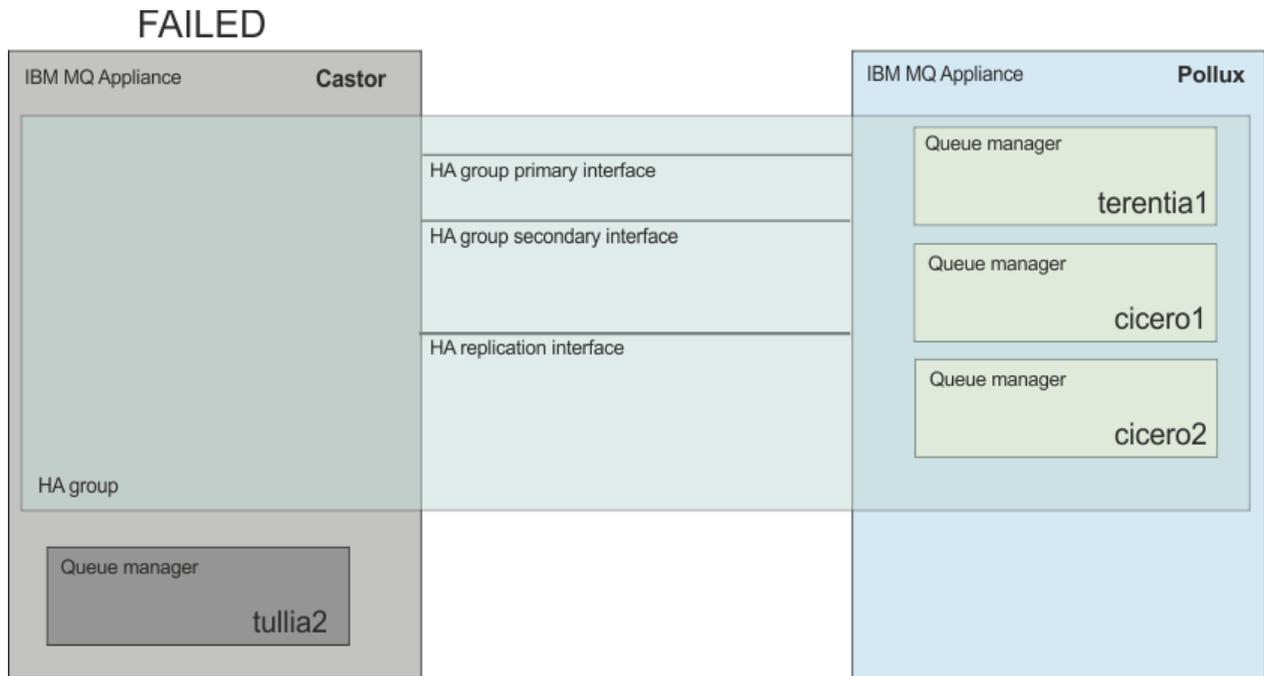


Figure 2. Example HA group after appliance failure

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](#)

Related concepts

[“Planning a high availability system” on page 33](#)

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

[“Configuring high availability” on page 267](#)

You can configure a pair of appliances to provide a high availability (HA) solution.

[“Operating in a high availability environment” on page 373](#)

If you have configured a high availability system, there are certain maintenance functions that you might need to complete.

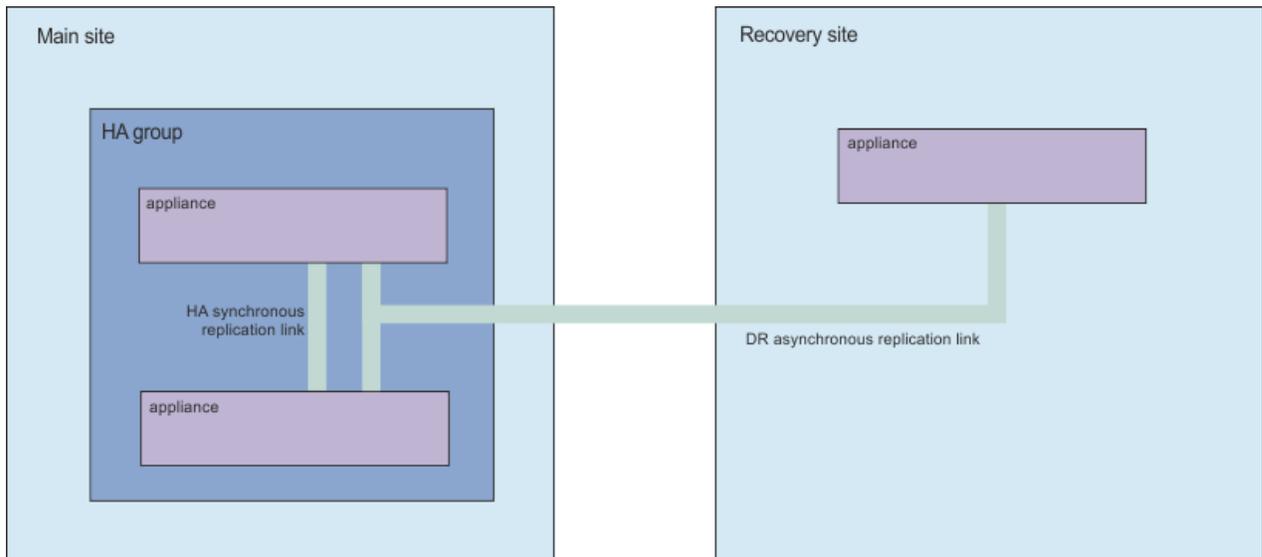
Disaster recovery for a high availability configuration

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.

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 18.

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” on page 577](#).

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, 2020. 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 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>.

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:

Command	Comment
crtmqcvx	Use of exits is not supported on the appliance. See “Exits and services on the IBM MQ Appliance” on page 31.
crtmqenv	Use the mqcli command line environment “setmqvar (set mq variable)” on page 623 command.
dltmqinst	See “Upgrading and downgrading firmware” on page 173.

Table 3. Unsupported commands (continued)

Command	Comment
dmpmqaut	Replaced by IBM MQ Console (see “Working with authority records” on page 329) or authority record configuration using runmqsc (for example, SET AUTHREC) see “runmqsc (run MQSC commands)” on page 606 .
dmpmqmsg	Use from an external system, connecting as a client.
dspmqaout	Replaced by IBM MQ Console (see “Working with authority records” on page 329) or authority record configuration using runmqsc (for example, SET AUTHREC) see “runmqsc (run MQSC commands)” on page 606 .
dspmqcsv	External control of some queue manager components is not supported on the IBM MQ Appliance.
dspmqls	IBM MQ file access is not supported on the appliance.
dspmqinst	See “Upgrading and downgrading firmware” on page 173
dspmqspl	Replaced by the SET/DISPLAY POLICY configuration using runmqsc, see “runmqsc (run MQSC commands)” on page 606 , see “Configuring IBM MQ Advanced Message Security” on page 265 .
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.
endmqweb	You cannot start or stop the mqweb server manually on the 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.
rcmqobj	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 amqsdq 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 318 , or use runmqsc, see “runmqsc (run MQSC commands)” on page 606 .
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 329) or authority record configuration using runmqsc (for example, SET AUTHREC) see “runmqsc (run MQSC commands)” on page 606 .

Table 3. Unsupported commands (continued)

Command	Comment
setmqenv	Use the mqcli command line environment “ setmqvar (set mq variable) ” on page 623 command.
setmqinst	See “ Upgrading and downgrading firmware ” on page 173
setmqm	See “ Upgrading and downgrading firmware ” on page 173.
setmqspl	Replaced by the SET/DISPLAY POLICY configuration using <code>runmqsc</code> , see “ runmqsc (run MQSC commands) ” on page 606, see “ Configuring IBM MQ Advanced Message Security ” on page 265.
setmqprd	See “ Upgrading and downgrading firmware ” on page 173
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.
strmqweb	You cannot start or stop the mqweb server manually on the 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 PCF and MQSC commands, see [Comparing command sets](#) in the IBM MQ documentation.

Unsupported queue manager parameters

The appliance does not support the queue manager security parameter **CONNAUTH CHCKLOCL**.

Supported commands

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

Table 4. Supported IBM MQ commands

Command	Comments
“ crtmqm (create queue manager) ” on page 578	<p>Create queue manager. The following parameters are not supported:</p> <ul style="list-style-type: none"> • -lc • -ll • -ld • -g • -md • -oa group • -q • -si • -ss • -z <p>The following parameters have been added:</p> <ul style="list-style-type: none"> • -fs <i>FileSize</i> • -sx
“ dlmqm (delete queue manager) ” on page 581	<p>Delete queue manager. The following parameter is not supported:</p> <ul style="list-style-type: none"> • -z

Table 4. Supported IBM MQ commands (continued)

Command	Comments
“dmpmqcfg (dump queue manager configuration)” on page 582	Dump queue manager configuration
“dspmq (display queue managers)” on page 585	<p>Display queue managers. The following parameters are not supported:</p> <ul style="list-style-type: none"> • -o installation • -o standby • -x <p>The following parameters have been added:</p> <ul style="list-style-type: none"> • -o ha • -o dr
“dspmqrte (display route information)” on page 586	Display route
“dspmqtrn (display incomplete transactions)” on page 594	Display transactions
“dspmqver (display version information)” on page 595	<p>Display version and build information. The following parameter is not supported:</p> <ul style="list-style-type: none"> • -i <p>The -p parameter supports only the values 1, 64, and 128.</p> <p>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.</p>
dspmqweb	Display information about the configuration of the mqweb server. The mqweb server is used to support the IBM MQ Console and administrative REST API.
“endmqm (end queue manager)” on page 597	<p>End queue manager. The following parameters are not supported:</p> <ul style="list-style-type: none"> • -s • -x • -z
“endmqtrc (end trace)” on page 599	End trace
“mqrc” on page 600	IBM MQ return code
“rcrmqobj (re-create object)” on page 601	Generate a client channel definition table (CCDT)
“rsvmqtrn (resolve transactions)” on page 602	Resolve transaction

Table 4. Supported IBM MQ commands (continued)

Command	Comments
“runmqras” on page 603	Run diagnostics collection. The following parameters are not supported: <ul style="list-style-type: none"> • -outputdir • -zipfile • -workdirectory
“runmqsc (run MQSC commands)” on page 606	Run MQSC commands. The following parameters are not supported: <ul style="list-style-type: none"> • -n • -c
“runswchl” on page 608	Switch cluster channel
“strmqm (start queue manager)” on page 610	Start queue manager. The following parameters are not supported: <ul style="list-style-type: none"> • -x • -z • -a • -I
“strmqtrc (Start trace)” on page 613	Start trace
“setmqvar (set mq variable)” on page 623	Add or remove an environment variable for the appliance or for a specified queue manager
setmqweb	Add or remove an mqweb server configuration property.

New commands

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

Table 5. Appliance commands

Command	Description
“dspmqerr (display mq error log)” on page 688	Display the IBM MQ error log files.
“crthagrp (create HA group)” on page 669	Create a high availability (HA) group of appliances.
“dsphagr (display HA group status)” on page 670	Display the status of the appliances in the high availability (HA) group.
“makehaprimary (make HA appliance primary)” on page 673	Specifies that an appliance is the 'winner' when resolving a partitioned situation in the high availability group.
“prepareha (prepare appliance for HA group)” on page 674	Prepare an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.
“sethagr (set HA group)” on page 675	Pause and resume an appliance in a high availability group.
“crtldrprimary” on page 680	Augment an existing queue manager to become the primary queue manager in a disaster recovery configuration.

Table 5. Appliance commands (continued)

Command	Description
“crtldrsecondary” on page 682	Create a secondary version of a queue manager on the recovery appliance in a disaster recovery configuration.
“makedrprimary” on page 684	Switch a disaster recovery queue manager to have the primary role in the disaster recovery configuration.
“makedrsecondary” on page 685	Prevent a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.
“dltdrprimary” on page 682	Remove a queue manager currently in the primary role from DR control.
“dltdrsecondary” on page 683	Remove a queue manager currently in the secondary role from DR control and delete it.
“dspmqini (display mq ini)” on page 617	Display attributes from the <code>qm.ini</code> or <code>mqat.ini</code> file of a specified queue manager.
“dspmqvar (display mq variables)” on page 619	Display environment variables set for a specified queue manager.
“setmqini (set mq ini)” on page 620	Add or remove an attribute from the <code>qm.ini</code> file of a specified queue manager. Set a value for an attribute in the <code>mqat.ini</code> file.
“addcert” on page 629	Add the public part of a certificate to the keystore of a specific queue manager.
“createcert (create certificate)” on page 630	Create a self-signed certificate for a queue manager.
“createcertrequest (create certificate request)” on page 632	Create a certificate request for a queue manager.
“deletecert (delete certificate)” on page 634	Delete a certificate from the keystore of a specific queue manager.
“deletecertrequest (delete certificate request)” on page 634	Delete a certificate request that was previously issued from a specific queue manager.
“detailcert (detail certificate)” on page 635	Show detailed information about a certificate for a specific queue manager.
“detailcertrequest (detail certificate request)” on page 636	Show detailed information about a certificate request for a specific queue manager.
“keybackup” on page 637	Back up the queue manager key repository to a file.
“keyrestore” on page 638	Restore a key repository
“listcert (list certificate)” on page 639	List the certificates that are held in the keystore of a specific queue manager.
“listcertrequest (list certificate request)” on page 640	List the certificate requests that are outstanding in the keystore of a specific queue manager.
“receivecert (receive certificate)” on page 641	Receive a certificate signed by a Certificate Authority (CA) as the result of a previous request.

Table 5. Appliance commands (continued)

Command	Description
“recreatecertrequest (re-create certificate request)” on page 642	Re-create a certificate request for a specific queue manager.
“usercreate (create user)” on page 624	Create user IDs for messaging users.
“userdelete (delete user)” on page 625	Delete a messaging user.
“usermodify (modify user)” on page 625	Modify user IDs for messaging users.
“userlist (list users)” on page 626	List the messaging users.
“groupcreate (create user group)” on page 626	Add user groups for messaging users.
“groupdelete (delete user group)” on page 627	Delete user groups for messaging users.
“grouplist (list user groups)” on page 627	List groups for messaging users.
“userbackup (back up users)” on page 627	Back up messaging users.
“userrestore (restore user)” on page 628	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 318](#)).

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 354](#)).

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. This 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 16](#) 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 synchronized 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 238).
- 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.
 - You can specify a floating IP address for individual HA queue managers. Applications can use the floating IP address to connect to a queue manager regardless of which appliance it is actually running on. You specify an interface name to connect on when you create the floating IP address (for example eth22). That interface must be a physical interface configured with a static IP address on both appliances. You can also specify an aggregate interface that exists with the same name on each appliance. The name of the aggregate interface must obey certain rules to be used for this purpose, see [“link-aggregation”](#) on page 803.
- 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 **V 9.1.0** (advanced users can, if required, use one of the appliance's other links, or an aggregated link, for replication).
 - 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”](#) on page 35.

The following diagram shows an example HA configuration:

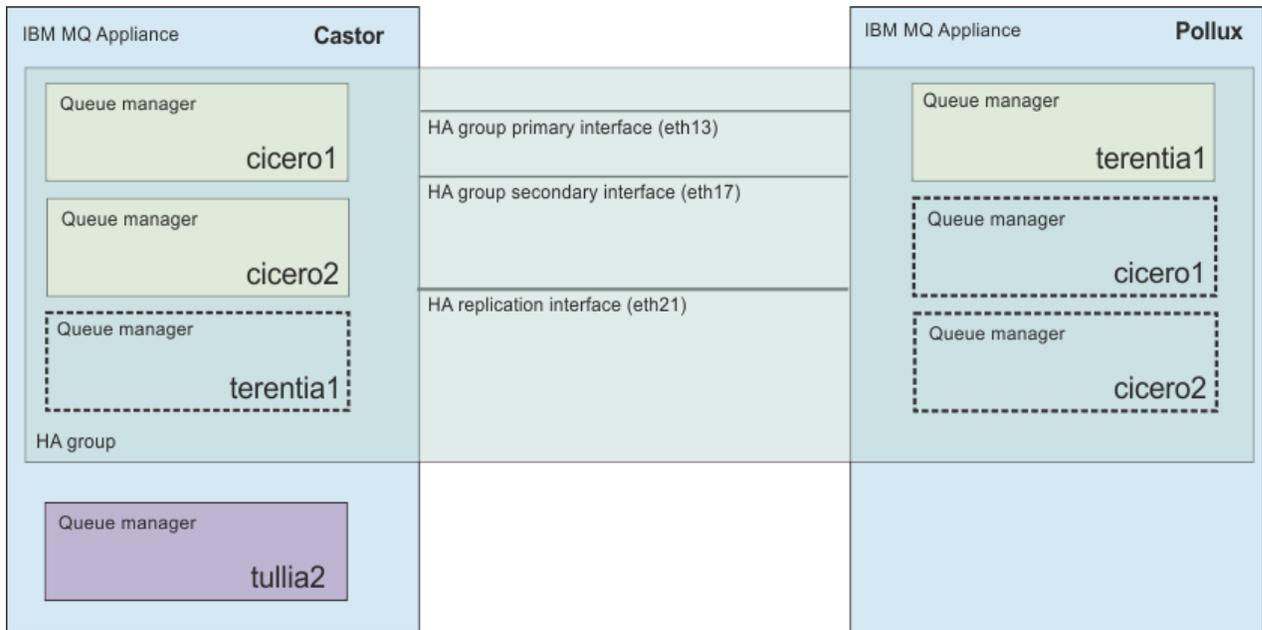


Figure 3. Example HA group

Related concepts

[“High availability” on page 16](#)

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.

[“Configuring high availability” on page 267](#)

You can configure a pair of appliances to provide a high availability (HA) solution.

[“Operating in a high availability environment” on page 373](#)

If you have configured a high availability system, there are certain maintenance functions that you might need to complete.

Network requirements for high availability

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

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 (anything longer is not supported), 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 **tracert** 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 207](#). 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 **tracert** 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 207](#). Log in as an administrative user.
2. Type the following command to test your connection:

```
tracert 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 **tracert** command with the IP addresses of the eth17 and eth21 interfaces of the other appliance in the high availability pair.

The **tracert** 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.

Related reference

[“ping” on page 714](#)

This command determines whether the network can reach a remote target.

[“tracert” on page 718](#)

This command traces the network path to a target host.

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 19](#) 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 18).
 - 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.
- Security
 - The link that is used to replicate queue manager data between the appliances is not subject to any secure encryption at the appliance level. As it is likely these connections will be wide area network connections that leave your secure enterprise network, it is important to make appropriate arrangements to encrypt these connections externally to the IBM MQ Appliance, for example, by using a hardware or software Virtual LAN product.

Related concepts

[“Disaster recovery”](#) on page 19

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.

[“Configuring disaster recovery”](#) on page 288

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.

Hardware limitations for disaster recovery

The following limitations apply to appliances configured for disaster recovery.

Network requirements for disaster recovery

- 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).

Related tasks

“Configuring the hardware for disaster recovery” on page 288

You must connect the appliances together via a network before you can configure the IBM MQ Appliance for disaster recovery (DR).

Planning network connections

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

The network connections are different on an M2001 appliance to an M2002 appliance.

M2002 Network connections on an M2002 appliance

The appliance has a number of network connections.

The network connections are situated on the front of an M2002 appliance. The location of the network connections is shown in the following illustration.

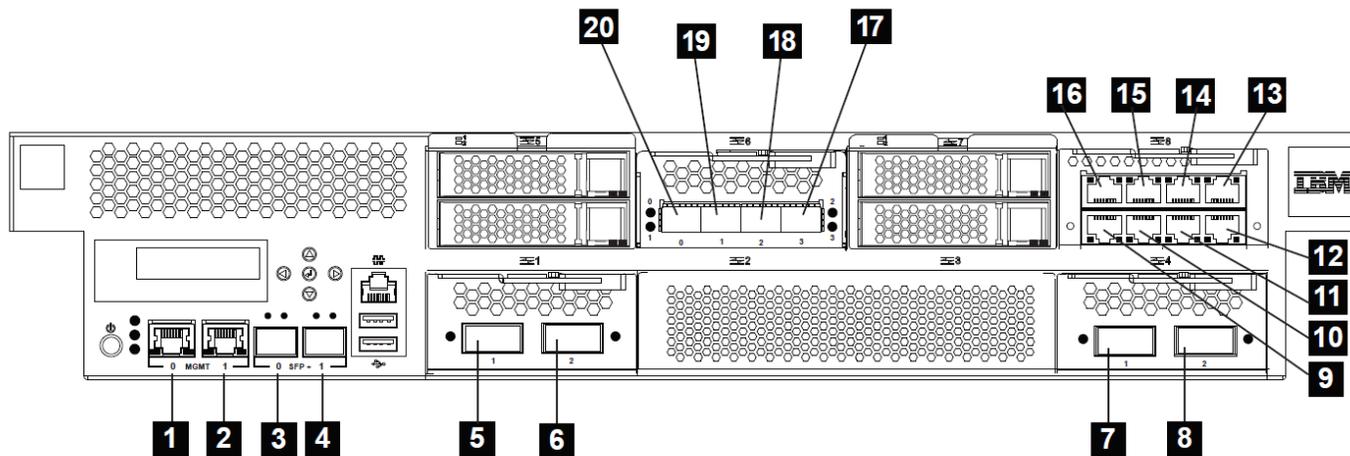


Figure 4. Network ports on an M2002 appliance

- 1** Management Ethernet port mgt0
- 2** Management Ethernet port mgt1
- 3** 10 Gb ethernet port, eth0
- 4** 10 Gb ethernet port, eth1
- 5** 40 Gb Ethernet module, eth30
- 6** 40 Gb Ethernet module, eth31
- 7** 40 Gb Ethernet module, eth32
- 8** 40 Gb Ethernet module, eth33
- 9** 1 Gb Ethernet module, eth10
- 10** 1 Gb Ethernet module, eth11

11

1 Gb Ethernet module, eth12

12

1 Gb Ethernet module, eth13

13

1 Gb Ethernet module, eth17

14

1 Gb Ethernet module, eth16

15

1 Gb Ethernet module, eth15

16

1 Gb Ethernet module, eth14

17

10 Gb Ethernet module, eth23

18

10 Gb Ethernet module, eth22

19

10 Gb Ethernet module, eth21

20

10 Gb Ethernet module, eth20

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 217).

Aggregate linking and configuration for VLANs is supported for connections. (VLAN is not supported where the connections are used in a high availability or disaster recovery configuration.)

Each 40 Gb Network module provides a total of 40Gb of bandwidth that is shared between the two physical ports on that module. You can create an aggregate interface using one physical connection to each network module (for example, eth30 and eth32) to boost total bandwidth and the availability characteristics of the link.

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 can use mgt0 or mgt1 for your IPMI connection. You can dedicate the mgt connection to IPMI, or you can share the mgt 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 mgt 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 the mgt port 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 (advanced users can define a different interface for replication, if required, see [“Configuring custom HA replication interfaces”](#) on page 269)

Ideally, you should have a dedicated subnet for each of the HA interfaces. The following table shows a set up for two appliances in an HA configuration, named *quirinus* and *romulus*, using IPv4 addresses and Class C private networks for each connection:

Interface	quirinus	romulus
eth13	192.168.13.149	192.168.13.150
eth17	192.168.17.149	192.168.17.150
eth21	192.168.21.149	192.168.21.150

Disaster recovery

If you are implementing a disaster recovery (DR) configuration, note that eth20 is dedicated to replication for DR (advanced users can define a different interface for replication, if required, see [“Configuring custom DR replication interfaces”](#) on page 290). The eth20 interface on each appliance in the DR configuration can be in a different subnet, although the subnets should be dedicated to DR replication.

If one of the appliances in a DR configuration also belongs to an HA configuration, the eth20 interfaces of the HA group must be in the same subnet, but the eth20 interface on the DR appliance can be in a different subnet.

Disaster recovery on high availability systems

If you are implementing DR on an HA system, then eth20 is used for DR replication by both HA appliances, and the recovery appliance (advanced users can define a different interface for replication, if required, see [“Configuring custom DR replication interfaces”](#) on page 290).

- 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 defined for eth20). 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 282). 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 823.

Table 7.

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 monitoring	log-int: 10.161.121.5
eth13	192.168.121.5/29	Used for HA Primary	hap-int: 192.168.121.5
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 M2002 appliance has four 10 Gb connections (eth0, eth1, eth22, eth23) and four 40 Gb connections (eth30, eth31, eth32, eth33) that are not by default used for HA or DR.

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. In this example, custom 40 Gb interfaces have been used for HA replication, and DR replication in place of the default eth20 and eth21 replication interfaces. A separate interface (aggregated for availability) is reserved for monitoring data such as writing to syslog targets, or SNMP infrastructure.

Table 8. Example network configuration for M2002 appliance

Interface	Used for
mgt0	Web UI and IPMI
mgt1	Command line access (SSH)
eth0	Not used
eth1	Not used

Table 8. Example network configuration for M2002 appliance (continued)

Interface	Used for
eth10	Not used
eth11	Not used
eth12	Link aggregated 1 Gb interface for monitoring
eth13	HA primary group interface
eth14	Not used
eth15	Not used
eth16	Link aggregated 1 Gb interface for monitoring
eth17	HA group alternative interface
eth20	not used
eth21	not used
eth22	Link aggregated 10 Gb data interface
eth23	Link aggregated 10 Gb data interface
eth30	Aggregated with eth32 for custom HA replication
eth31	Aggregated with eth33 for custom DR replication
eth32	see eth30
eth33	see eth31

M 2001 Network connections on an M2001 appliance

The appliance has a number of network connections.

The network connections are situated on the front of an M2001 appliance. The location of the network connections is shown in the following illustration.

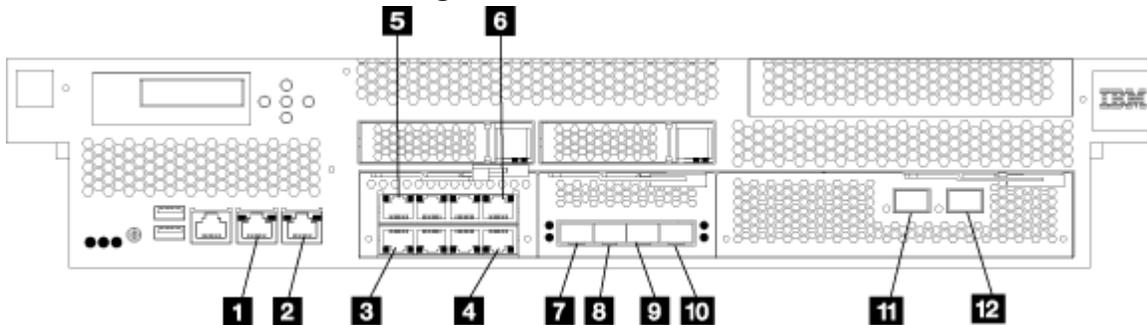


Figure 5. Network ports on an M2001 appliance

- 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 Gb Ethernet module, eth20
- 8** 10 Gb Ethernet module, eth21
- 9** 10 Gb Ethernet module, eth22 (M2001 appliance only)
- 10** 10 Gb 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 217).

Aggregate linking and configuration for VLANs is supported for connections (VLAN is not supported for high availability or disaster recovery configurations).

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 (advanced users can define a different interface for replication, if required, see [“Configuring custom HA replication interfaces”](#) on page 269)

Ideally, you should have a dedicated subnet for each of the HA interfaces. The following table shows a set up for two appliances in an HA configuration, named *quirinus* and *romulus*, using IPv4 addresses and Class C private networks for each connection:

Interface	quirinus	romulus
eth13	192.168.13.149	192.168.13.150
eth17	192.168.17.149	192.168.17.150
eth21	192.168.21.149	192.168.21.150

Disaster recovery

If you are implementing a disaster recovery (DR) configuration, note that eth20 is dedicated to replication for DR (advanced users can define a different interface for replication, if required, see [“Configuring custom DR replication interfaces”](#) on page 290). The eth20 interface on each appliance in the DR configuration can be in a different subnet, although the subnets should be dedicated to DR replication.

If one of the appliances in a DR configuration also belongs to an HA configuration, the eth20 interfaces of the HA group must be in the same subnet, but the eth20 interface on the DR appliance can be in a different 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 (advanced users can define a different interface for replication, if required, see [“Configuring custom DR replication interfaces”](#) on page 290).

- 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 defined for eth20). 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 282). 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 823.

Table 10.

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 monitoring	log-int: 10.161.121.5
eth13	192.168.121.5/29	Used for HA Primary	hap-int: 192.168.121.5
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. A separate interface (aggregated for availability) is reserved for monitoring data such as writing to syslog targets, or SNMP infrastructure.

Table 11. 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 monitoring

Table 11. Example network configuration for M2001 appliance (continued)

Interface	Used for
eth13	HA primary group interface
eth14	Not used
eth15	Not used
eth16	Link aggregated 1 Gb interface for monitoring
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

Network configuration guidance

You can configure your own network connections on the IBM MQ Appliance using this guidance to help.

One of the advantages of the appliance is that all the administration tasks can be carried out by a single appliance administrator. This guidance helps you to set up networking on the appliance even if you are not yourself a networking expert.

When you install firmware on your appliance for the first time, you can set up one management interface and a default gateway as part of the running the installation wizard, which is enough to set up connectivity with the outside world. You might well require a more sophisticated network configuration, however, and as the appliance has many network interfaces, you need to plan before you configure additional interfaces. You need to consider the following points:

- What is the topology of the network that you are connecting to?
 - Do you have a dedicated management subnet?
 - Do your brokers need to connect to multiple subnets?
- What is the motivation for configuring multiple appliance connections? (performance, redundancy, or security are possible reasons)

TCP/IP network routing function

In simple terms, when the appliance needs to route to a host, it assesses available connections in this order:

1. Is there an existing interface on the same subnet as the target host?
2. Are there static routes defined to that specific host?
3. Are there static routes defined to that host's subnet?
4. Is there a default gateway defined?

You can use the **show route** command to display the information currently available to the appliance in making these decisions. The **show route** command shows the appliance routing table. The table includes static and default routes from appliance interface configurations.

The aim in configuring your appliance is to avoid any ambiguity or uncertainty when routing to a host. Ambiguity can cause problems for some network operations, for example, when pinging an appliance, you might see no response if the return path is different to the request path. Such ambiguity can also interfere with the high availability and disaster recovery functionality of the appliance.

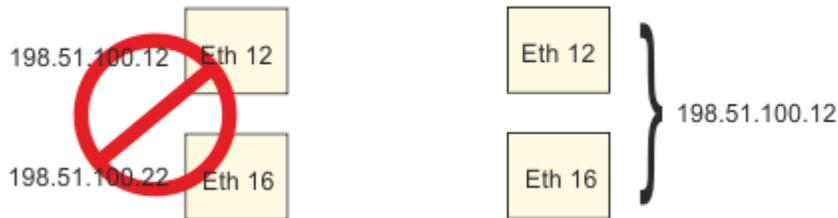
Best Practice guidelines

Although much of your configuration will be dictated by the structure of the network that you are connecting to, and your priorities in terms of performance, redundancy, and security, you can follow these guidelines to help avoid ambiguity and uncertainty.

Avoid having multiple IP addresses on the same subnet allocated to appliance network interfaces

If you are planning to do this to provide redundancy, consider using link aggregation. You can aggregate several of the appliance interfaces together, using a single IP address to access them (see [“Link aggregation interfaces”](#) on page 222).

*do not configure interfaces
on same subnet ...*

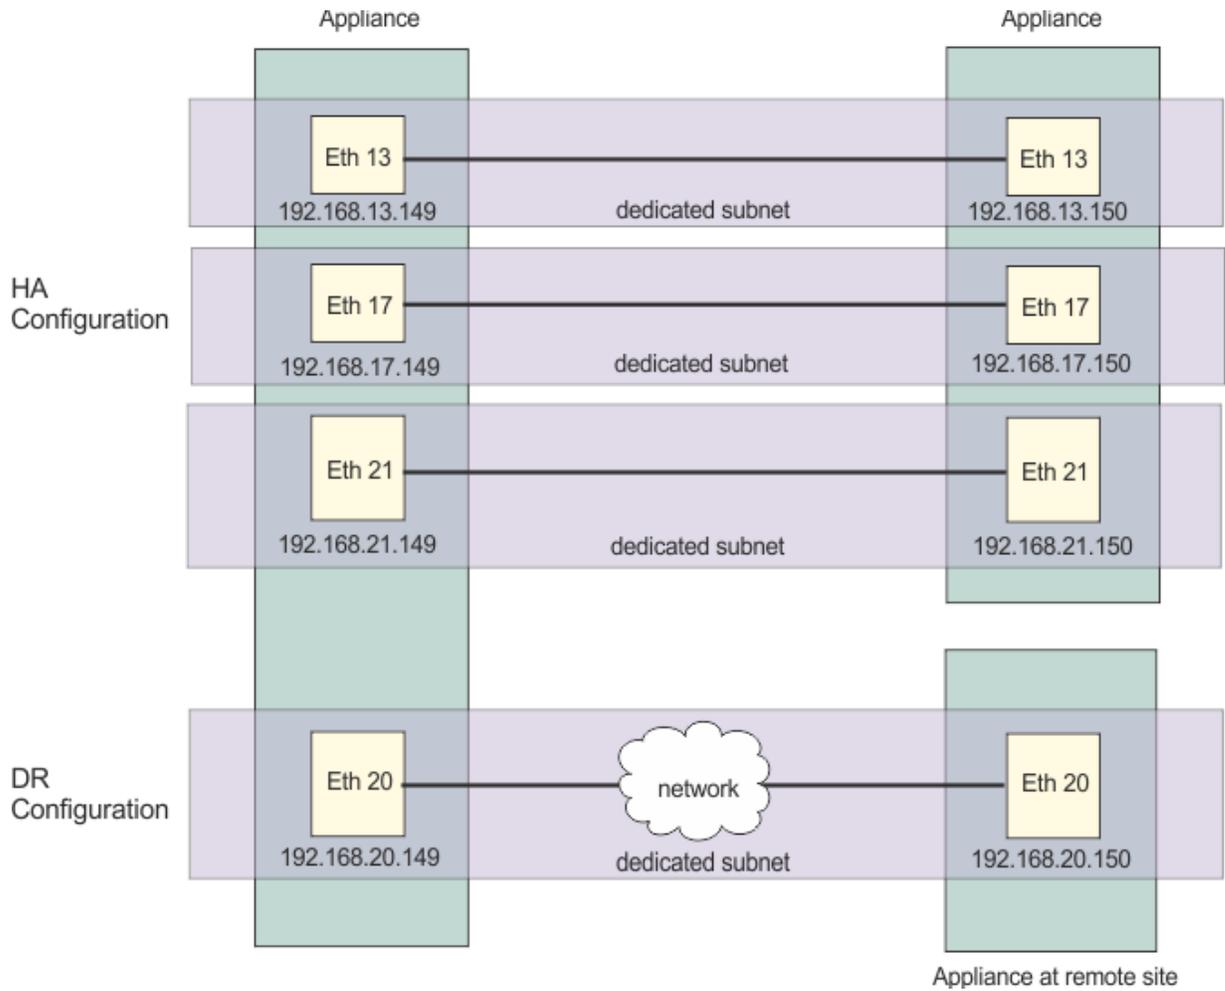


...use link aggregation instead

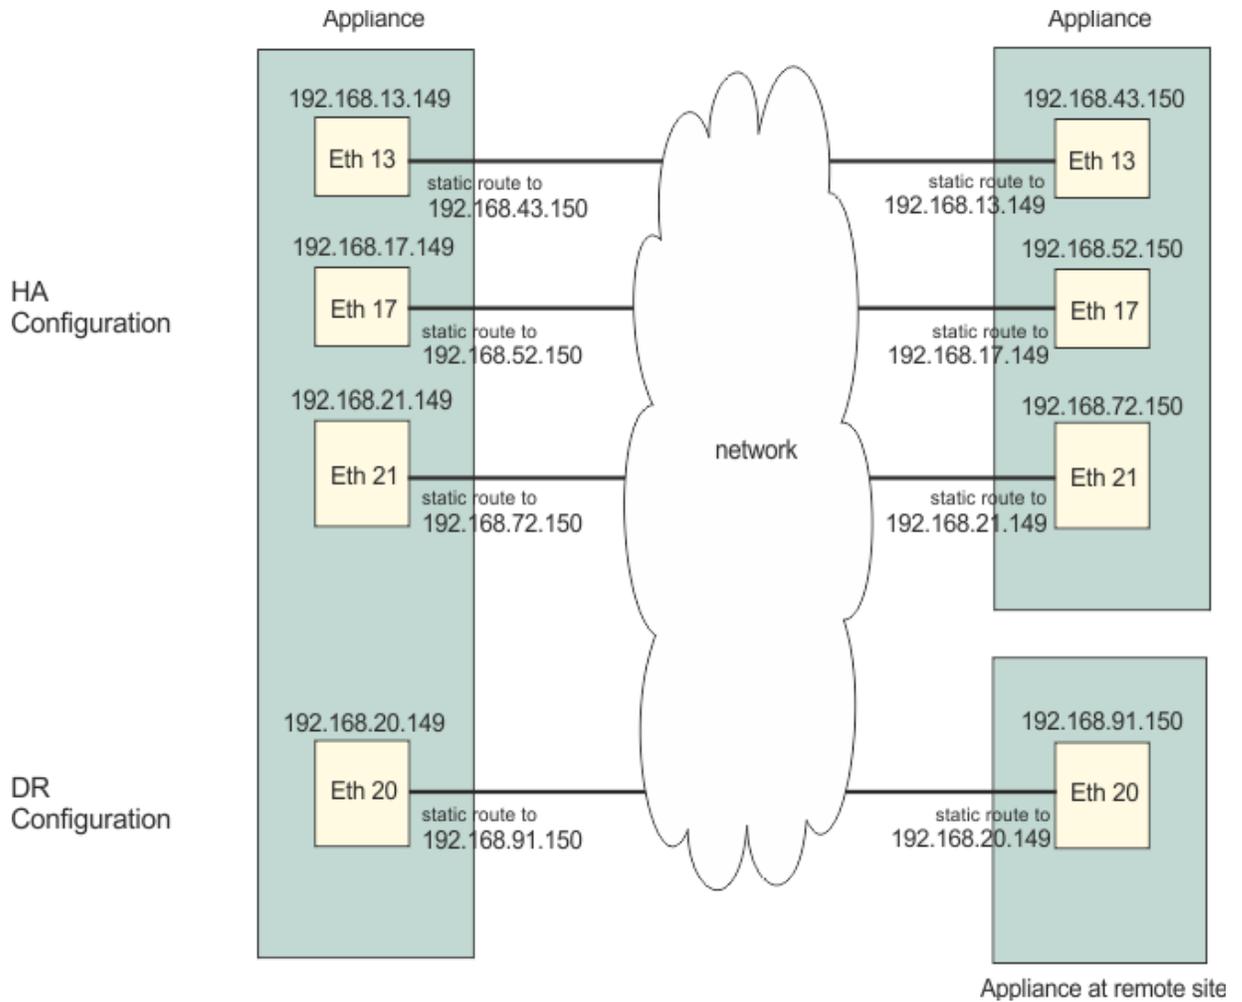
Configure HA and DR connections into separate, dedicated subnets, or define static routes

Put direct HA and DR connections into separate, dedicated subnets. Giving each direct connection its own subnet will completely remove any potential issues for clashes. Such connections do not need gateways or routers of any kind, since all traffic on these direct connections will be peer-to-peer within that subnet.

If you are not using direct cable connections for your HA or DR interfaces, you should still use discrete dedicated subnets for each connection (this is most likely to be true for your DR connection, which would usually be at a different site rather than physically nearby as for HA systems).

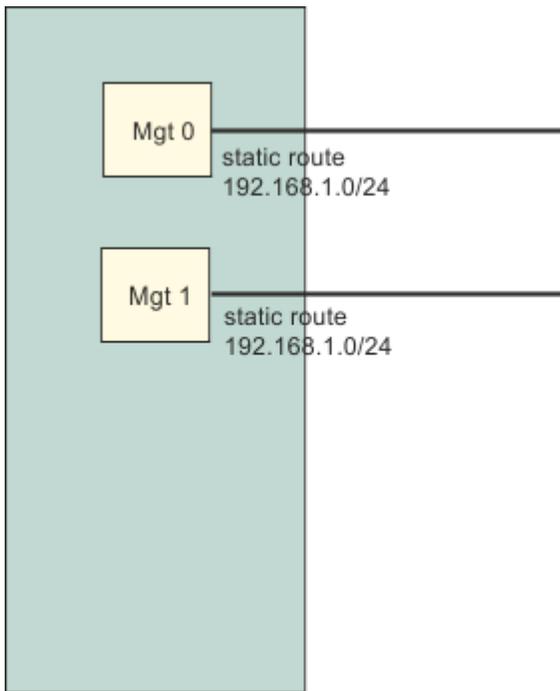


If you cannot configure dedicated subnets, define static routes for your HA and DR connections.



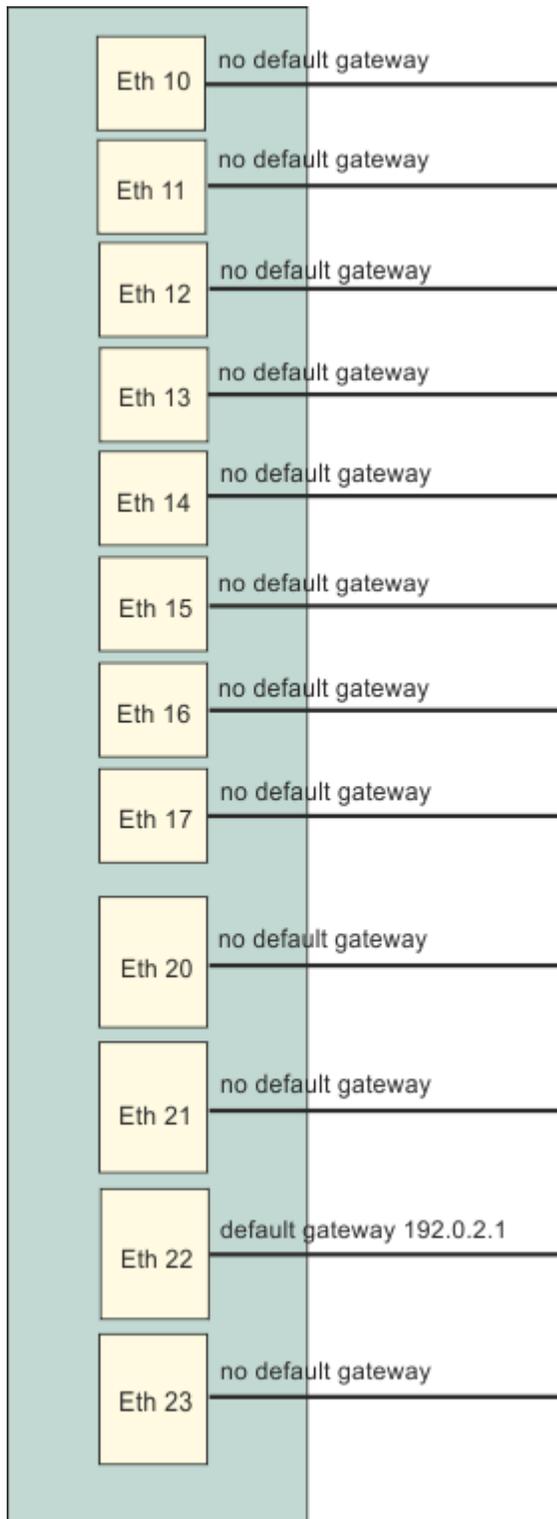
Consider defining separate static routes to hosts or subnets for specific MQ and appliance management traffic

For example, if you know that all of your management traffic should be coming to and from 192.168 (private network) addresses, define static routes on mgt0 and/or mgt1 to ensure that traffic with these systems takes a known route and does not interfere with other (for example, application) traffic.



Define only one default gateway and on one interface

Avoid unpredictable routing by defining only one default gateway, and define it on a single interface. Define the default gateway on one of the interfaces you expect outgoing MQ connections to use, as this makes it easy for queue managers to route outwards to any IP that does not have a more specific route defined.



► M 2000 M 2001 Planning SAN storage

You can configure an IBM MQ queue manager to use a storage area network (SAN) for queue manager data.

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances. Support for SAN is now deprecated and there will be no further functional enhancements to this feature.

By default, queue managers use the built-in RAID storage on the appliance. You can specify that a non-high availability queue manager stores all its data on SAN storage if you require enhanced capacity, speed, and resilience. The appliance has a fibre channel interface for connection to the SAN.

You must first configure your appliance to use the SAN. You can then specify that a queue manager uses SAN-storage when you create it on the appliance.

You must configure your SAN to provide one or more dedicated partitions for appliance use. Each queue manager uses a separate storage end point, conceptually a separate disk or device. SAN storage does not provide support for shared storage (for example, multi-instance queue managers).

The appliance supports access to a switched SAN fabric accessed by using fibre channel host bus adapters. You must configure your SAN so that each queue manager uses a separate, dedicated LUN. Ensure that your storage network is zoned so that only the appliance using a particular volume can access the LUN in normal operations. In a disaster recovery scenario, it might be appropriate for multiple appliances to have access to a single LUN, although only one appliance has the volume active (enabled) at any given time.

Ensure administrative access to the data stored on SAN volumes is controlled and audited appropriately.

Queue manager data is not encrypted by the appliance, so you must take steps to secure your SAN storage independently. Volumes configured for appliance use are used to store sensitive information, including certificates and password files relating to the queue manager.

Related concepts

[“Configuring SAN storage” on page 300](#)

You can configure the IBM MQ Appliance so that queue managers can use SAN storage.

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.
- You must allocate disk storage to each queue manager from your available storage when you create the queue manager. An HA or DR queue manager requires twice the storage of a stand-alone queue manager.

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](#)

Related concepts

[“Appliance comparison: M2002 and M2001” on page 12](#)

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

Security planning

Security planning involves thinking about administrative access to the appliance itself, and thinking about implementing security for IBM MQ messaging.

User security

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.

Administrative access of users to both the appliance and to IBM MQ is controlled by role based management. Using role based management you can finely control exactly what resources users have access to. Users can be authenticated by using an LDAP repository, a local or remote XML file, or can be configured as local users. User permissions can be defined in an XML file or a local user group. The appliance always has the privileged administrative account `admin`, that you cannot configure. There are some appliance operations that can only be performed by the `admin` user.

Messaging users are defined using IBM MQ commands. You can use role based management delegate the IBM MQ authority checks from an appliance user to a matching messaging user. This means that you can effectively have a single user ID to cover both functions.

For more information about role based management, see [“Role based management” on page 439](#).

For more information about messaging users, see [“Administering messaging users” on page 347](#).

For more information about how users are authorized to use different resources, see [“User authorization, credential mapping, and access profiles” on page 423](#).

For examples of how to set up different users on your appliance to access different resources, see [“Configuring user access to the IBM MQ Console and the CLI” on page 254](#).

Link level security for IBM MQ

The appliance supports the Transport Layer Security (TLS) protocol to provide link level security for message channels and MQI channels. The appliance supports the same levels of TLS as IBM MQ, although the implementation details are different.

For more information about TLS in IBM MQ, see [SSL and TLS security protocols](#) in the IBM MQ documentation.

For more information about how the appliance implements TLS, see [“TLS certificate management” on page 479](#).

Message level security for IBM MQ

If you require a higher level of security for sensitive data flowing through the IBM MQ, you can implement IBM MQ Advanced Message Security on the appliance. Under AMS, certificates are distributed to IBM MQ clients, and the clients then encrypt and decrypt data at the application. Use of AMS guarantees that message data has not been modified between when it is originally placed on a queue and when it is retrieved. In addition, AMS verifies that a sender of message data is authorized to place signed messages on a target queue.

MCA interceptor

If you are unable to configure some or all of your IBM MQ clients to encrypt and decrypt data at the application, you can use an AMS MCA interceptor to give some of the benefits of AMS without making changes to the client. Using an MCA interceptor, data is encrypted/decrypted as it enters or leaves the queue manager or queue manager network by the channel itself, and therefore only users/applications with access to the appropriate certificates can decrypt data as stored internally to the queue manager. Use this mechanism where managing the encryption at the endpoint is not an option.

Only client applications that cannot manage encryption should connect on interceptor-enabled channels.

Use of the MCA interceptor can give you increased protection against users without physical access to the appliance, including any appliance administrative users who do not have access to the `mqcli` command. The intercept can thus give increased protection against 'privilege escalation' style attacks from other users.

Use of the MCA interceptor should not, however, be regarded as equivalent to full end-to-end encryption using client-side certificates, or (on platforms where available) full disk encryption. The certificates that are used to encrypt or sign messages are stored on the appliance disks, alongside the actual queue manager data. This storage means that administrators who have access to the `mqcli` command on the appliance can request that a copy of this certificate store is exported from the system. Also, anyone with physical access to the appliance could potentially remove the disks and access the certificate store by installing the disks in another system. You should bear these limitations in mind when planning to use MCA intercepts.

For more information on IBM MQ Advanced Message Security, see [IBM MQ Advanced Message Security](#) in the IBM MQ documentation.

For more information on implementing IBM MQ Advanced Message Security on the appliance, see [“Configuring IBM MQ Advanced Message Security” on page 265](#).

For more information on implementing MCA intercepts on the appliance, see [“Configuring MCA interception” on page 266](#).

IBM MQ and IBM MQ Appliance on premises considerations for GDPR readiness

For PID(s):

- 5724-H72 IBM MQ
- 5655-AV9 IBM MQ Advanced for z/OS®
- 5655-AV1 IBM MQ Advanced for z/OS Value Unit Edition
- 5655-AM9 IBM MQ Advanced Message Security for z/OS
- 5725-S14 IBM MQ Appliance M2000
- 5725-Z09 IBM MQ Appliance M2001
- 5737-H47 IBM MQ Appliance M2002
- 5655-MQ9 IBM MQ for z/OS
- 5655-VU9 IBM MQ for z/OS Value Unit Edition
- 5655-MF9 IBM MQ Managed File Transfer for z/OS
- 5655-ADV IBM WebSphere® MQ Advanced for z/OS
- 5655-AMS IBM WebSphere MQ Advanced Message Security for z/OS
- 5724-A39 IBM WebSphere MQ for HP NonStop Server
- 5655-W97 IBM WebSphere MQ for z/OS
- 5655-VU8 IBM WebSphere MQ for z/OS Value Unit Edition
- 5655-VUE IBM WebSphere MQ for z/OS Value Unit Edition
- 5655-MFT IBM WebSphere MQ Managed File Transfer for z/OS

Notice:

This document is intended to help you in your preparations for GDPR readiness. It provides information about features of IBM MQ that you can configure, and aspects of the product's use, that you should consider to help your organization with GDPR readiness. This information is not an exhaustive list, due to

the many ways that clients can choose and configure features, and the large variety of ways that the product can be used in itself and with third-party applications and systems.

Clients are responsible for ensuring their own compliance with various laws and regulations, including the European Union General Data Protection Regulation. Clients are solely responsible for obtaining advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulations that may affect the clients' business and any actions the clients may need to take to comply with such laws and regulations.

The products, services, and other capabilities described herein are not suitable for all client situations and may have restricted availability. IBM does not provide legal, accounting, or auditing advice or represent or warrant that its services or products will ensure that clients are in compliance with any law or regulation.

Table of Contents

1. [GDPR](#)
2. [Product Configuration for GDPR](#)
3. [Data Life Cycle](#)
4. [Data Collection](#)
5. [Data Storage](#)
6. [Data Access](#)
7. [Data Processing](#)
8. [Data Deletion](#)
9. [Data Monitoring](#)
10. [Capability for Restricting Use of Personal Data](#)

GDPR

General Data Protection Regulation (GDPR) has been adopted by the European Union ("EU") and applies from May 25, 2018.

Why is GDPR important?

GDPR establishes a stronger data protection regulatory framework for processing of personal data of individuals. GDPR brings:

- New and enhanced rights for individuals
- Widened definition of personal data
- New obligations for processors
- Potential for significant financial penalties for non-compliance
- Compulsory data breach notification

Read more about GDPR:

- [EU GDPR Information Portal](#)
- [ibm.com/GDPR website](http://ibm.com/GDPR)

Product Configuration - considerations for GDPR Readiness

The following sections provide considerations for configuring IBM MQ to help your organization with GDPR readiness.

Data Life Cycle

IBM MQ is a transactional message oriented middleware product that enables applications to asynchronously exchange application provided data. IBM MQ supports a range of messaging APIs,

protocols and bridges for the purpose of connecting applications. As such, IBM MQ may be used to exchange many forms of data, some of which could potentially be subject to GDPR. There are several third-party products with which IBM MQ might exchange data. Some of these are IBM-owned, but many others are provided by other technology suppliers. The [Software Product Compatibility Reports website](#) provides lists of the associated software. For considerations regarding the GDPR readiness of a third-party product, you should consult that product's documentation. IBM MQ administrators control the way in which IBM MQ interacts with data passing through it, by the definition of queues, topics and subscriptions.

What types of data flow through IBM MQ?

As IBM MQ provides asynchronous messaging service for application data, there is no one definitive answer to this question because use cases vary through application deployment. Application message data is persisted in queue files (page sets or the Coupling Facility on z/OS), logs and archives and the message may itself contain data that is governed by GDPR. Application provided message data may also be included in files collected for problem determination purposes such as error logs, trace files and FFSTs. On z/OS application provided message data may also be included in address space or Coupling Facility dumps.

The following are some typical examples of personal data that may be exchanged using IBM MQ:

- Employees of the customer (for example; IBM MQ might be used to connect the customer's payroll or HR systems)
- The customer's own clients' personal data (for example; IBM MQ might be used by a customer to exchange data between applications that relates to their clients, such as taking sales leads and storing data inside their CRM system).
- The customer's own clients' sensitive personal data (for example; IBM MQ might be used within industry contexts that require personal data to be exchanged, such as HL7-based healthcare records when integrating clinical applications).

In addition to application provided message data, IBM MQ processes the following types of data:

- Authentication Credentials (such as username and passwords, API keys, etc.)
- Technically Identifiable Personal Information (such as device IDs, usage based identifiers, IP address, etc. - when linked to an individual)

Personal data used for online contact with IBM

IBM MQ clients can submit online comments/feedback/requests to contact IBM about IBM MQ subjects in a variety of ways, primarily:

- Public comments area on pages in the [IBM MQ area on IBM Developer](#)
- Public comments area on pages of [IBM MQ product documentation in IBM Knowledge Center](#)
- Public comments in the [IBM Support Forums](#)
- Public comments in the [IBM RFE Community on IBM Developer](#)

Typically, only the client name and email address are used, to enable personal replies for the subject of the contact, and the use of personal data conforms to the [IBM Online Privacy Statement](#).

Data Collection

IBM MQ can be used to collect personal data. When assessing your use of IBM MQ and your needs to meet with the demands of GDPR, you should consider the types of personal data which in your circumstances are passing through IBM MQ. You may wish to consider aspects such as:

- How does data arrive into your queue managers? (Across which protocols? Is the data encrypted? Is the data signed?)
- How is data sent out from your queue managers? (Across which protocols? Is the data encrypted? Is the data signed?)
- How is data stored as it passes through a queue manager? (Any messaging application has the potential to write message data to stateful media, even if a message is non-persistent. Are you aware of how

messaging features could potentially expose aspects of the application message data passing through the product?)

- How are credentials collected and stored where needed by IBM MQ to access third-party applications?

IBM MQ may need to communicate with other systems and services which require authentication, for example LDAP. Where needed, authentication data (userids, passwords) is configured and stored by IBM MQ for its use in such communications. Wherever possible, you should avoid using personal credentials for IBM MQ authentication. Consider the protection of the storage used for authentication data. (See Data Storage below.)

Data Storage

When message data travels through queue managers, IBM MQ will persist (perhaps multiple copies of) that data directly to stateful media. IBM MQ users may want to consider securing message data whilst it is at rest.

The following items highlight areas where IBM MQ persists application provided data, which users may wish to consider when ensuring compliance with GDPR.

- Application Message Queues:

IBM MQ provides message queues to allow asynchronous data exchange between applications. Non-persistent and persistent messages stored on a queue are written to stateful media.

- File Transfer Agent Queues:

IBM MQ Managed File Transfer utilizes message queues to co-ordinate the reliable transfer of file data, files that contain personal data and records of transfers are stored on these queues.

- Transmission Queues:

To transfer messages reliably between queue managers, messages are stored temporarily on transmission queues.

- Dead-Letter Queues:

There are some circumstances in which messages cannot be put to a destination queue and are stored on a dead-letter queue, if such a queue is configured on the queue manager.

- Backout Queues:

JMS and XMS messaging interfaces provide a capability that allows poison messages to be moved to a backout queue after a number of backouts have occurred to allow other valid messages to be processed.

- AMS Error Queue:

IBM MQ Advanced Message Security will move messages that don't comply with a security policy to the SYSTEM.PROTECTION.ERROR.QUEUE error queue in a similar way to dead-letter queuing.

- Retained Publications:

IBM MQ provides a retained publication feature to allow subscribing applications to recall a previous publication.

Read more:

- [Logging: Making sure that messages are not lost](#)
- [MFT Agent queue settings](#)
- [Using the dead-letter queue](#)
- [Handling poison messages in IBM MQ classes for JMS](#)
- [AMS error handling](#)
- [Retained publications](#)

The following items highlight areas where IBM MQ may indirectly persist application provided data which users may also wish to consider when ensuring compliance with GDPR.

- Trace route messaging:

IBM MQ provides trace route capabilities, which record the route a message takes between applications. The event messages generated may include technically identifiable personal information such as IP addresses.

- Application activity trace:

IBM MQ provides application activity trace, which record the messaging API activities of applications and channels, application activity trace can record the contents of application provided message data to event messages.

- Service trace:

IBM MQ provides service trace features, which records the internal code paths through which message data flows. As part of these features, IBM MQ can record the contents of application provided message data to trace files stored on disk.

- Queue manager events:

IBM MQ can generate event messages that could include personal data, such as authority, command and configuration events.

Read more:

- [Trace-route messaging](#)
- [Using trace](#)
- [Event monitoring](#)
- [Queue manager events](#)

To protect access to copies of the application provided message data consider the following actions:

- Restrict privileged user access to IBM MQ data in the filesystem, for example restricting user membership of the 'mqm' group on UNIX platforms.
- Restrict application access to IBM MQ data via dedicated queues and access control. Where appropriate avoid unnecessary sharing of resources such as queues between applications and provide granular access control to queue and topic resources.
- Restrict access to replicated copies of IBM MQ data in high availability (HA) or disaster recovery (DR) configurations, and secure the connections used for replication.
- Use IBM MQ Advanced Message Security to provide end-to-end signing and/or encryption of message data.
- Use file- or volume-level encryption to protect the contents of the directory used to store trace logs.
- After uploading service trace to IBM, you can delete your service trace files and FFST data if you are concerned about the contents potentially containing personal data.

Read more:

- [Privileged users](#)
- [Planning file system support on Multiplatforms](#)

An IBM MQ administrator may configure a queue manager with credentials (username and password, API keys, etc.) for 3rd party services such as LDAP, Salesforce, etc. This data is generally stored in the queue manager data directory protected through file system permissions.

When an IBM MQ queue manager is created, the data directory is set up with group-based access control such that IBM MQ can read the configuration files and use the credentials to connect to these systems. IBM MQ administrators are considered privileged users and are members of this group so have read access to the files. Some files are obfuscated but they are not encrypted. For this reason, to fully protect access to credentials, you should consider the following actions:

- Restrict privileged user access to IBM MQ data, for example restricting membership of the 'mqm' group on UNIX platforms.
- Use file- or volume-level encryption to protect the contents of the queue manager data directory.

- Encrypt backups of the production configuration directory and store them with appropriate access controls.
- Consider providing audit trails for authentication failure, access control and configuration changes with security, command and configuration events.

Read more:

- [Securing IBM MQ](#)

Data Access

IBM MQ queue manager data can be accessed through the following product interfaces, some of which are designed for access through a remote connection, and others for access through a local connection.

- IBM MQ Console [Only Remote]
- IBM MQ Administrative REST API [Only Remote]
- IBM MQ Messaging REST API [Only Remote]
- MQI [Local and Remote]
- JMS [Local and Remote]
- XMS [Local and Remote]
- IBM MQ Telemetry (MQTT) [Only Remote]
- IBM MQ Light (AMQP) [Only Remote]
- IBM MQ IMS bridge [Only Local]
- IBM MQ CICS bridge [Only Local]
- IBM MQ MFT Protocol bridges [Only Remote]
- IBM MQ Connect:Direct bridges [Only Remote]
- IBM MQ Bridge to Salesforce [Only Remote]
- IBM MQ Bridge to Blockchain [Only Remote]
- IBM MQ MQAI [Local and Remote]
- IBM MQ PCF commands [Local and Remote]
- IBM MQ MQSC commands [Local and Remote]
- IBM MQ Explorer [Local and Remote]
- IBM MQ User Exits [Only Local]
- IBM MQ Internet Pass-Thru [Only Remote]
- IBM MQ Appliance Serial Console [Only Local]
- IBM MQ Appliance SSH [Only Remote]
- IBM MQ Appliance REST API [Only Remote]
- IBM MQ Appliance Web UI [Only Remote]

The interfaces are designed to allow users to make changes to an IBM MQ queue manager and messages stored on it. Administration and messaging operations are secured such that there are three stages involved when a request is made;

- Authentication
- Role mapping
- Authorization

Authentication:

If the message or administrative operation was requested from a local connection, the source of this connection is a running process on the same system. The user running the process must have passed any authentication steps provided by the operating system. The user name of the owner of the process from which the connection was made is asserted as the identity. For example, this could be the name of the

user running the shell from which an application has been started. The possible forms of authentication for local connections are:

1. Asserted user name (local OS)
2. Optional username and password (OS, LDAP or custom 3rd party repositories)

If the administrative action was requested from a remote connection, then communications with IBM MQ are made through a network interface. The following forms of identity may be presented for authentication via network connections;

1. Asserted user name (from remote OS)
2. Username and password (OS, LDAP or custom 3rd party repositories)
3. Source network address (such as IP address)
4. X.509 Digital Certificate (mutual SSL/TLS authentication)
5. Security tokens (such as LTPA2 token)
6. Other custom security (capability provided by 3rd party exits)
7. SSH keys

Role mapping:

In the role mapping stage, the credentials that were provided in the authentication stage may be mapped to an alternate user identifier. Provided the mapped user identifier is permitted to proceed (for example administrative users may be blocked by channel authentication rules), then the mapped user id is carried forward into the final stage when authorizing activities against IBM MQ resources.

Authorization:

IBM MQ provides the ability for different users to have different authorities against different messaging resources such as queues, topics and other queue manager objects.

Logging activity:

Some users of IBM MQ may need to create an audit record of access to MQ resources. Examples of desirable audit logs might include configuration changes that contain information about the change in addition to who requested it.

The following sources of information are available to implement this requirement:

1. An IBM MQ queue manager can be configured to produce command events when an admin command has been run successfully.
2. An IBM MQ queue manager can be configured to produce configuration events when a queue manager resource is created, altered or deleted.
3. An IBM MQ queue manager can be configured to produce an authority event when an authorization check fails for a resource.
4. Error messages indicating failed authorization checks are written to the queue manager error logs.
5. The IBM MQ Console will write audit messages to its logs when authentication, authorization checks fail or when queue managers are created, started, stopped or deleted.
6. The IBM MQ Appliance will write audit messages to its logs to record user log ins and system changes.

When considering these kind of solutions, IBM MQ users might want to give consideration to the following points:

- Event messages are non-persistent so when a queue manager restarts the information is lost. Any event monitors should be configured to constantly consume any available messages and transfer the content to persistent media.
- IBM MQ privileged users have sufficient privileges to disabled events, clear logs or delete queue managers.

For more information about securing access to IBM MQ data and providing an audit trail refer to the following topics:

- [IBM MQ security mechanisms](#)
- [Configuration events](#)
- [Command events](#)
- [Error logs](#)

Data Processing

Encryption using a Public Key Infrastructure:

You can secure network connections to IBM MQ by specifying that the connections use TLS, which can also provide mutual authentication of the initiating side of the connection.

Using the PKI security facilities that are provided by transport mechanisms is the first step towards securing data processing with IBM MQ. However, without enabling further security features, the behavior of a consuming application is to process all messages delivered to it without validating the origin of the message or whether it was altered whilst in transit.

Users of IBM MQ that are licensed to use Advanced Message Security (AMS) capabilities can control the way in which applications process personal data held in messages, through the definition and configuration of security policies. Security policies allow digital signing and/or encryption to be applied to message data between applications.

It is possible to use security policies to require and validate a digital signature when consuming messages to ensure messages are authentic. AMS encryption provides a method by which message data is converted from a readable form to an encoded version that can only be decoded by another application if it is the intended recipient or the message and has access to the correct decryption key.

For more information about using SSL and certificates to secure your network connections, refer to the following topics in the IBM MQ product documentation:

- [Configuring TLS security for IBM MQ](#)
- [AMS Overview](#)

Data Deletion

IBM MQ provides commands and user interface actions to delete data that has been supplied to the product. This means that users of IBM MQ can delete data that relates to particular individuals, should this be required.

- Areas of IBM MQ behavior to consider for complying with GDPR Client Data deletion
 - Delete message data stored on an application queue by:
 - Removing individual messages using messaging API or tooling or by using message expiry.
 - Specifying that messages are non-persistent, held on a queue where non-persistent message class is normal and restarting the queue manager.
 - Administratively clearing the queue.
 - Deleting the queue.
 - Delete retained publication data stored on a topic by:
 - Specifying that messages are non-persistent and restarting the queue manager.
 - Replacing the retained data with new data or by using message expiry.
 - Administratively clearing the topic string.
 - Delete data stored on a queue manager by deleting the whole queue manager and any replicated copies for high availability or disaster recovery.
 - Delete data stored by the Service trace commands by deleting the files in the trace directory.
 - Delete FFST data stored by deleting the files in the errors directory.
 - Delete address space and Coupling Facility dumps (on z/OS).

- Delete archive, backup or other copies of such data.
- Areas of IBM MQ behavior to consider for complying with GDPR Account Data deletion
 - You can delete account data and preferences stored by IBM MQ for connecting to queue managers and 3rd party services by deleting (including archive, backup or otherwise replicated copies thereof):
 - Queue manager authentication information objects that store credentials.
 - Queue manager authority records that reference user identifiers.
 - Queue manager channel authentication rules that map or block specific IP addresses, certificate DN's or user identifiers.
 - Credentials files used by IBM MQ Managed File Transfer Agent, Logger and MQ Explorer MFT Plugin for authentication with queue manager and file servers.
 - X.509 digital certificates that represent or contain information about an individual from keystores which may be used by SSL/TLS connections or IBM MQ Advanced Message Security (AMS).
 - Individual user accounts from IBM MQ Appliance, including reference to those accounts in system log files.
 - IBM MQ Explorer workspace metadata and Eclipse settings.
 - IBM MQ Explorer password store as specified in the [Password Preferences](#).
 - IBM MQ Console and mqweb server configuration files.
 - Salesforce connection data configuration files.
 - Blockchain connection data configuration files.
 - IBM MQ Internet Pass-Thru configuration files and keystores.

Read more:

- [Configuring the IBM MQ Bridge to Salesforce](#)
- [Configuring IBM MQ for use with blockchain](#)
- [MFT and IBM MQ connection authentication](#)
- [Mapping credentials for a file server by using the ProtocolBridgeCredentials.xml file](#)
- [Configuring IBM MQ Console users and roles](#)

Data Monitoring

IBM MQ provides a range of monitoring features that users can exploit to gain a better understanding of how applications and queue managers are performing.

IBM MQ also provides a number a features that help manage queue manager error logs.

Read more:

- [Monitoring your IBM MQ network](#)
- [Diagnostic message services](#)
- [QMErrorLog service](#)
- [IBM MQ Appliance monitoring and reporting](#)

Capability for Restricting Use of Personal Data

Using the facilities summarized in this document, IBM MQ enables an end-user to restrict usage of their personal data.

IBM MQ message queues should not be used as a permanent data store in the same way as a database, which is particularly true when handling application data that is subject to GDPR.

Unlike a database where data may be found through a search query, it can be difficult to find message data unless you know the queue, message and correlation identifiers of a message.

Provided messages containing an individual's data can be readily identified and located, it is possible using standard IBM MQ messaging features to access or modify message data.

M 2001 Installing the M2001 appliance

Plan for your installation, install, and verify the installation of the IBM MQ Appliance, model M2001.

Installation demonstration

View the video to see a demonstration of a IBM MQ Appliance

 [Installing an IBM MQ Appliance](#)

M 2001 Safety

Before you install this product, read the Safety Information.

Arabic

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الأمنية.

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 ασφαλείας (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.

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 pečítajte 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.

M 2001 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.

M 2001 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.

M 2001 Safety statements

Safety statements are available on the included CD-ROM.

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

M 2001 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.

M 2001 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)

M 2001 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)

M 2001 Labels

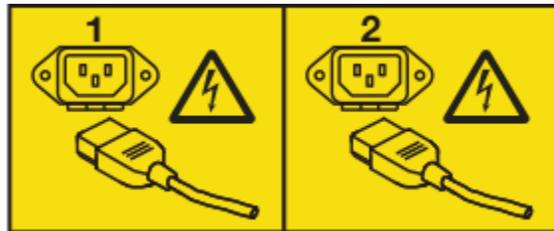
One or more of the following safety labels may apply to this product.



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)

M 2001 Introducing the IBM MQ Appliance M2001

The IBM MQ Appliance provides IBM MQ V9 on an easy-to-deploy appliance.

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

M 2001 Specifications and features

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

M 2001 Specifications

Hardware specifications for the appliance.

The following table summarizes the specifications for the chassis.

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)

M 2001 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.

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.

M 2001 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.

M 2001 Components on the front

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

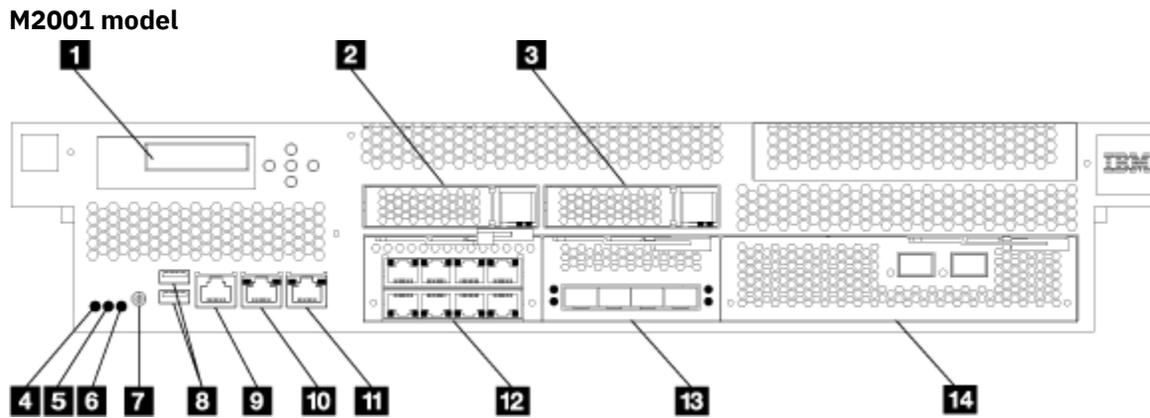


Figure 6. Controls, connectors, and status indicators on the front of the appliance.

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

- 1** LCD display.
- 2** 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

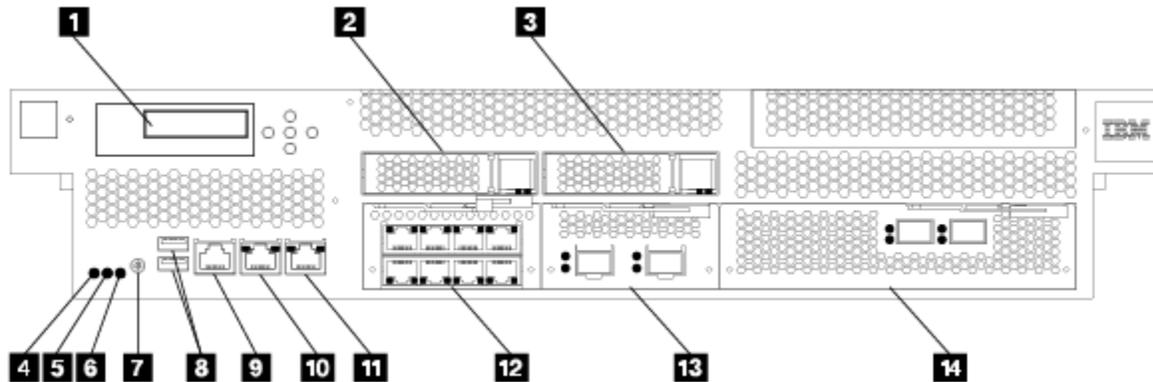


Figure 7. Controls, connectors, and status indicators on the front of the appliance.

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.

M 2001 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.

M 2001 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`.

M 2001 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.

Related tasks

[“Turning off the appliance” on page 105](#)

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

M 2001 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 terminal¹ to the appliance or to connect from a PC that is running terminal emulation software to the appliance.

M 2001 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 (`mg \dagger 0` and `mg \dagger 1`) 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.

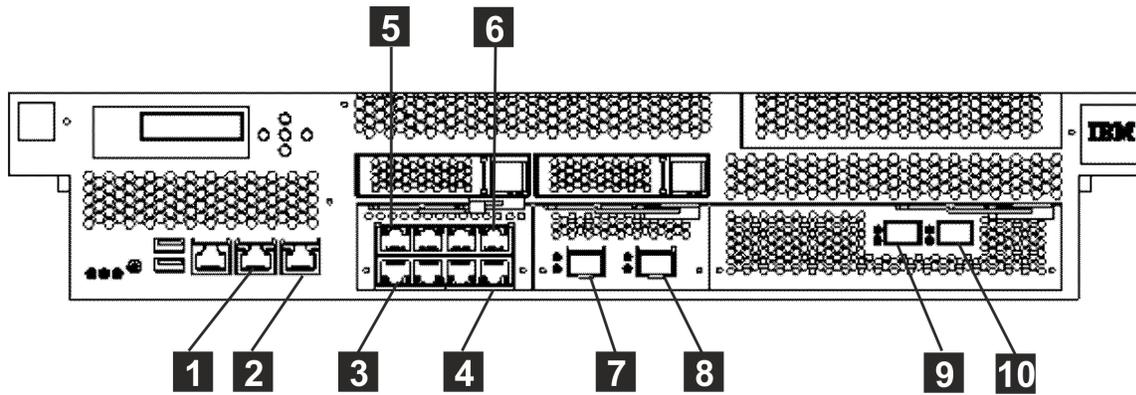


Figure 8. M2000 model network ports

- 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

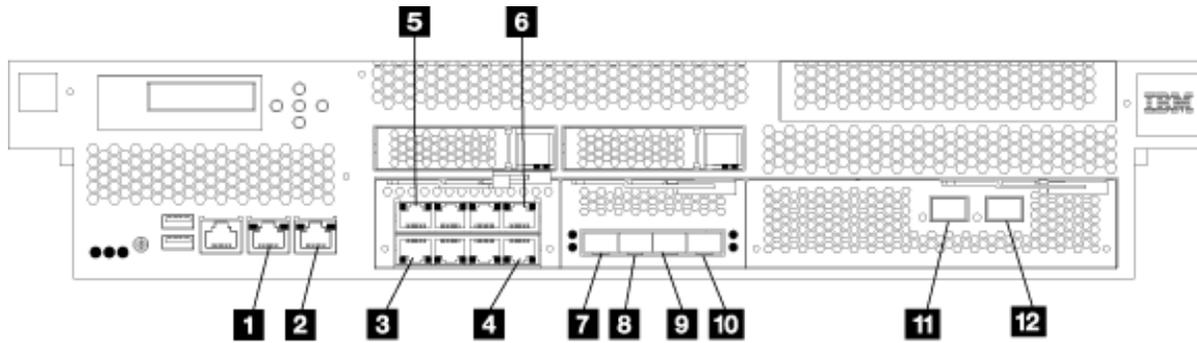


Figure 9. M2001 model network ports

- 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

M 2001 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.

M 2001 *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.

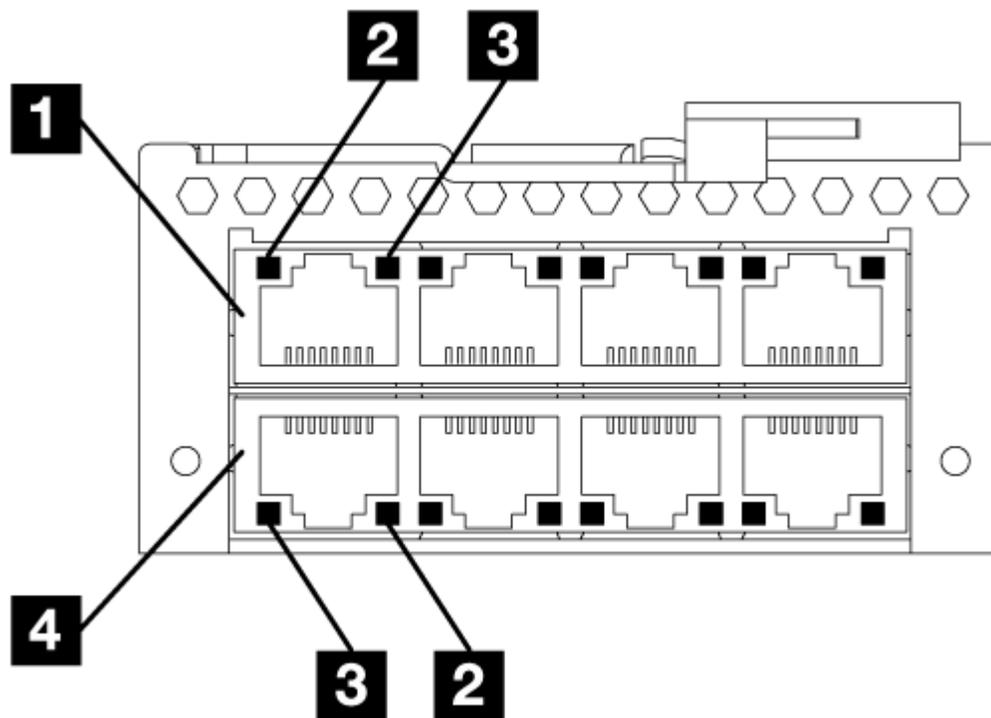


Figure 10. 1 Gb Ethernet module with eight ports for RJ45 interface

- 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.

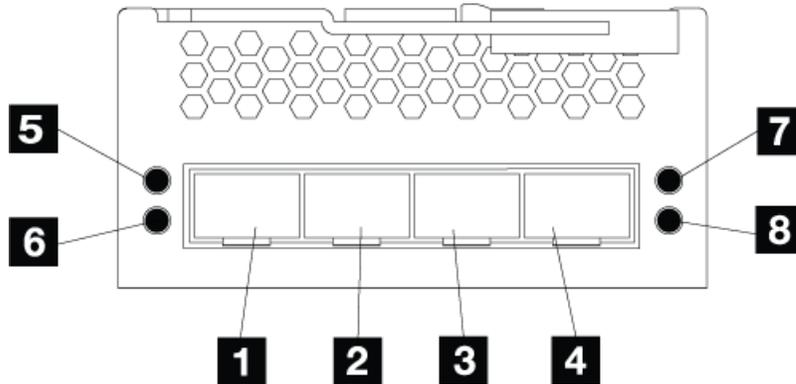


Figure 11. 10 Gb Ethernet module with two ports for SFP+ interface

- 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

M2001 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.

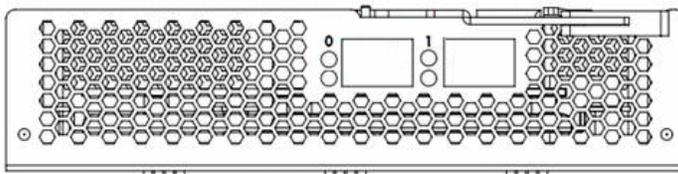


Figure 12. 16 Gb Fibre channel module

M 2001 Solid-state disk drive modules

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

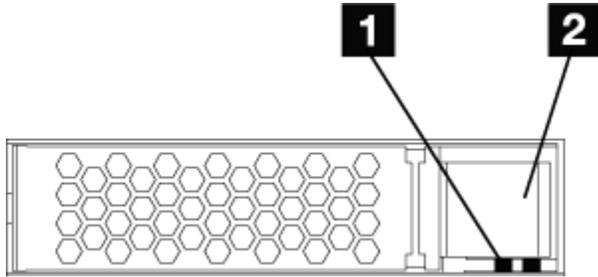


Figure 13. Solid-state disk drive module.

- 1** Solid-state disk drive activity LED.
- 2** Locking arm release latch.

M 2001 Hard disk drive modules

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

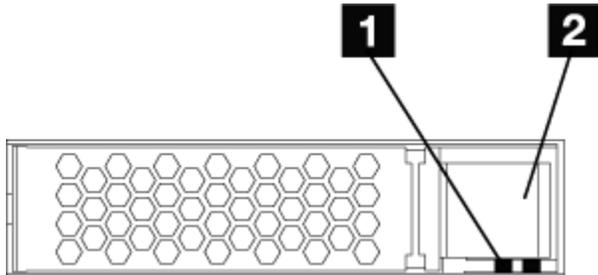


Figure 14. Hard disk drive module.

- 1** Hard disk drive activity LED.
- 2** Locking arm release latch.

M 2001 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.

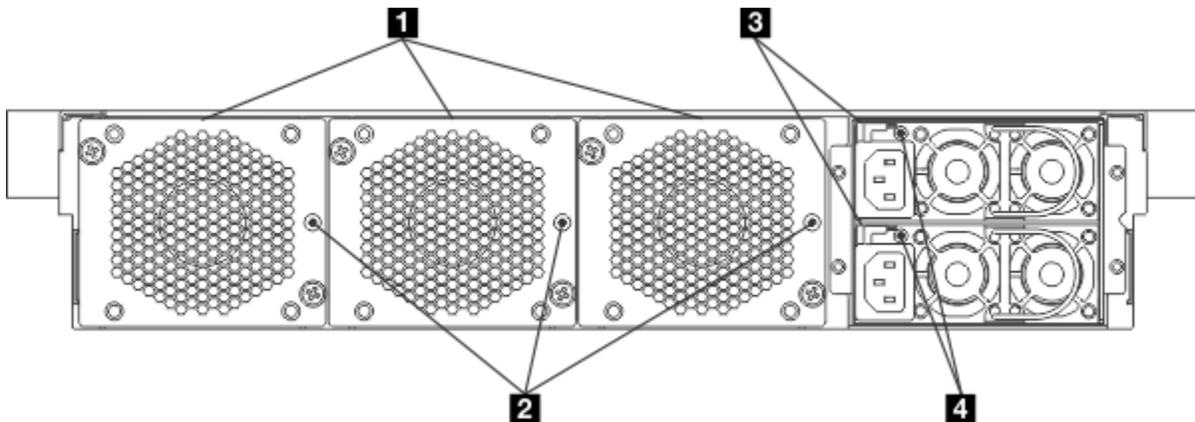


Figure 15. Rear view.

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.

M 2001 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.

M 2001 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)

Related tasks

“Turning off the appliance” on page 105

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

M 2001 Prepare for installation

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

M 2001 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)

M 2001 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.

M 2001 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.

M 2001 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.

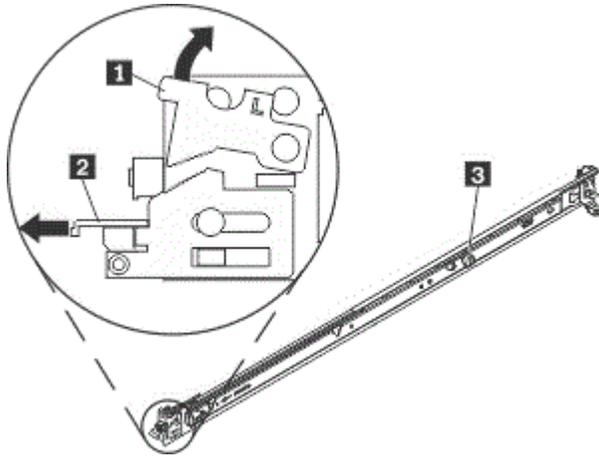


Figure 16. View of the left slide rail.

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.

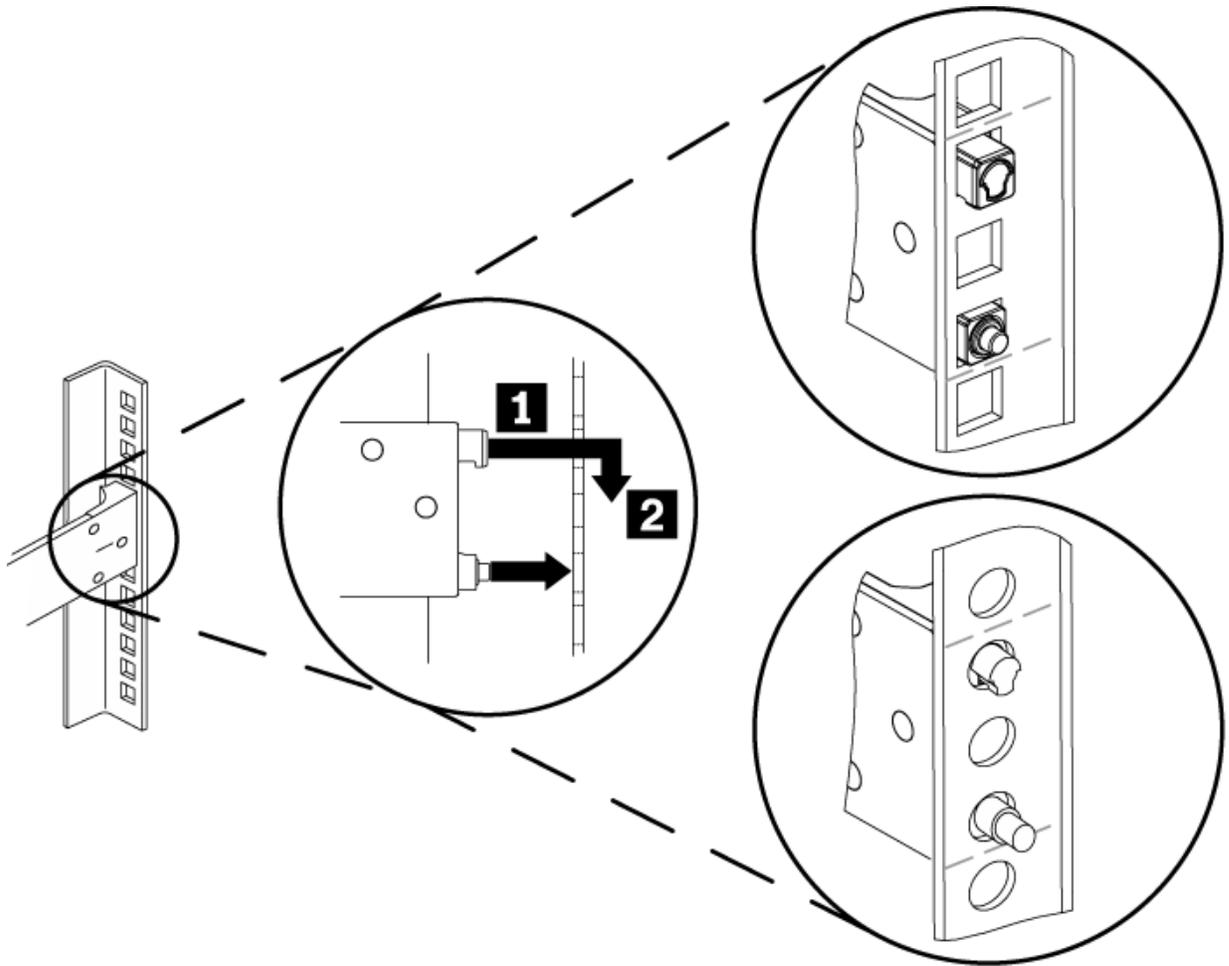


Figure 17. Install the rear end of the slide rails.

- 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.

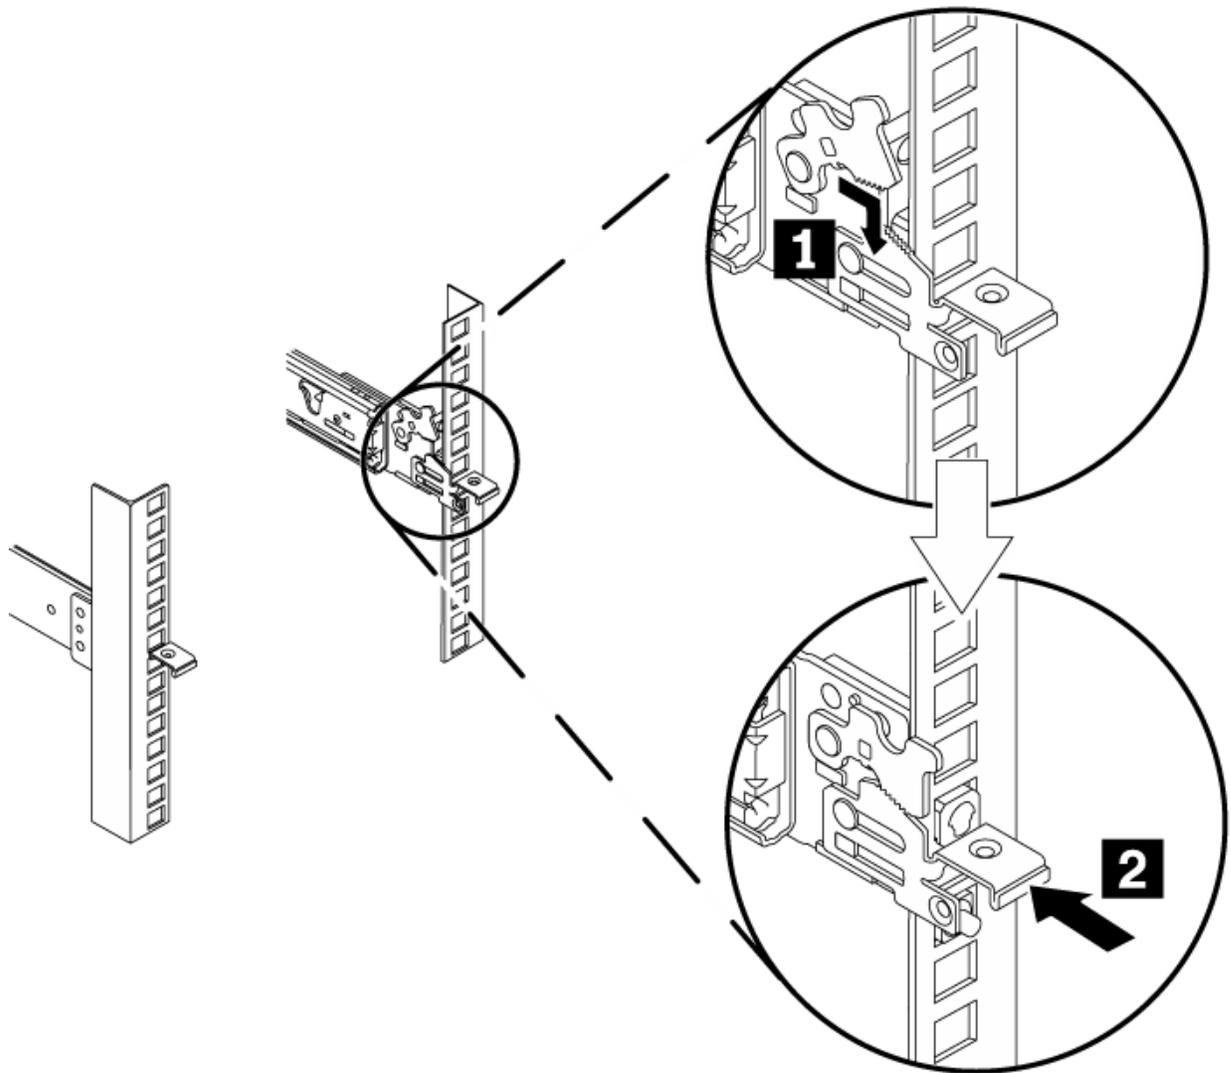


Figure 18. Install the front end of the slide rails.

- 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.

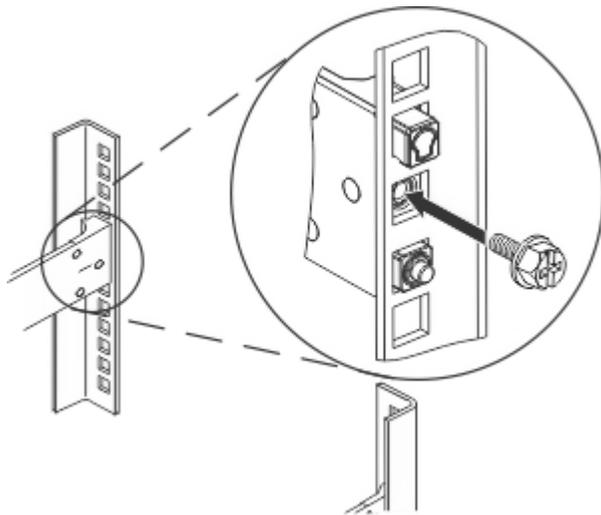


Figure 19. Securing the rails in the rack.

6. Repeat step 5 for the left rail.

► M2001 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.

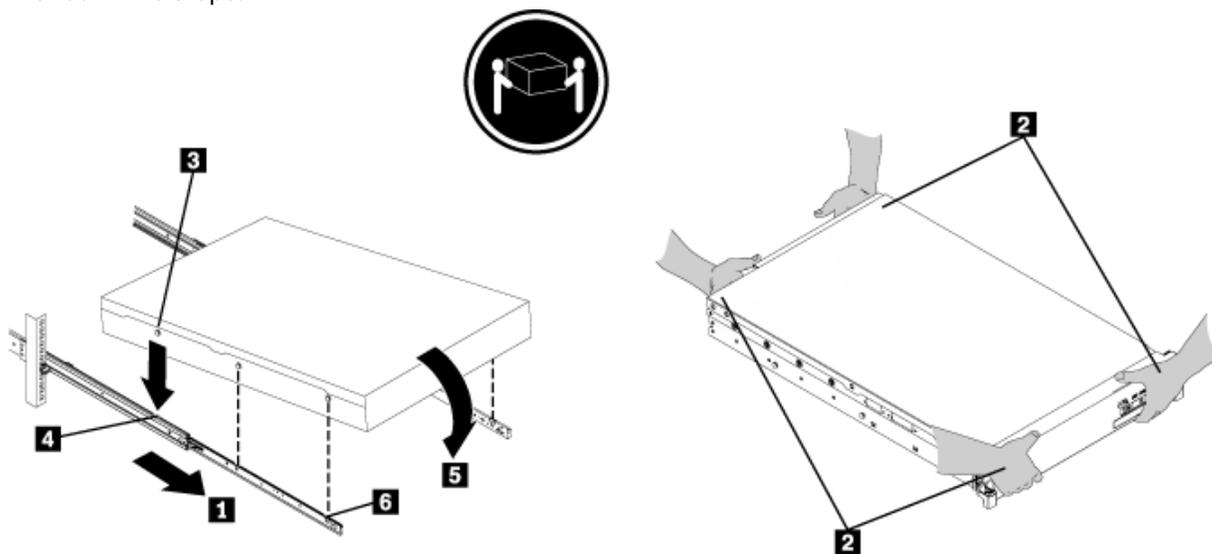


Figure 20. Securing the appliance in the rack



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.

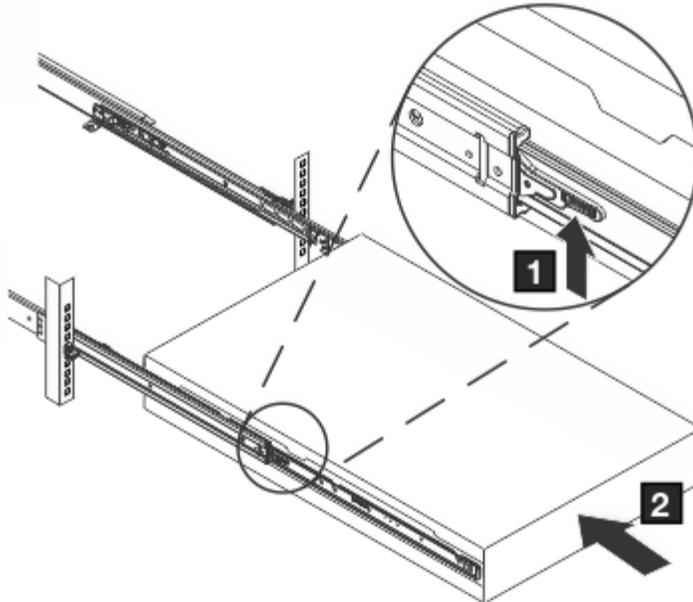


Figure 21. Sliding the appliance into the rack.

Procedure

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

M 2001 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.

M 2002 Connect the appliance to your network

Considerations to review before you connect the appliance to your 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: 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 those specified herein might result in hazardous radiation exposure. (C026)



CAUTION: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, 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.

The 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.

10Gb connectivity is provided through network modules that require an external transceiver (10Gb SFP+) appropriate to your chosen physical network. IBM provides short range (SR) optical transceivers with the appliance. These transceivers conform to the following specifications.

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

All optical transceivers provided by IBM are Class 1 Eye safe per requirements of IEC 60825-1 /CDRH.

A DAC (Direct attach copper) 10GB cable is also supplied with each appliance. This cable can be used in place of the supplied optical transceivers, for example, if directly connecting two appliances for high availability replication.

You might choose to replace the supplied transceivers in the 10GB ports with alternative parts, for example, to replace a short range optical transceiver with a long range optical transceiver. Any SFP+ transceiver (or DAC/AOC) following industry standards is compatible with the IBM MQ Appliance network modules. If a support problem is identified with an SFP+ transceiver not supplied by IBM, you might be asked to reproduce the problem on a component supplied by IBM. The 'link' and 'activity' LEDs can be used to confirm that the transceiver is operational.

The network modules in the IBM MQ Appliance are based on the following Intel adapters, and any transceivers known to be compatible with these adapters are therefore appropriate:

- 4-port SFP+ 10GbE NIM is equivalent to the Intel Ethernet Converged Network Adapter X710-DA4

Related tasks

[“Removing an SFP+ transceiver” on page 114](#)

[The procedure to remove a 10 Gb SFP+ transceiver.](#)

M 2001 **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 89](#) and [“Appliance modes” on page 89](#))
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.

M 2001 **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.

M 2001 **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.

M 2001 **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`.

M 2001 Firmware considerations

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

M 2001 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.

M 2001 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 is aimed at larger enterprise workloads.
- IBM MQ Appliance M2001B 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 appliance to an M2001B+ appliance, which has the same capacity as an M2001A appliance.

Upgrading from an M2001B to M2001B+ 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.

V 9.1.4 You can view which mode is currently operational on the appliance in the following ways:

- By viewing the LCD display panel on the front of the appliance.
- In the welcome banner when you start the command line interface.
- By using the **show system** command (reported as Capacity mode). See “[show](#)” on page 885.
- In the System Setting panel of the appliance Web UI (displayed as **Licensed capacity mode**).
- By using a REST API command to show system settings. See “[REST management interface](#)” on page 1000.

M 2001 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 88, and read the operation modes and password considerations for the `admin` account in “[Firmware considerations](#)” on page 89.

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, 8N1², and no flow control data transfer.

What to do next

See “[Procedure 2 of 3: Initializing the appliance](#)” on page 90 to define the base configuration such as changing the password for the `admin` account, creating privileged users, and configuring the web management interface.

M 2001 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 90 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.

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.

- 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 admin³.
3. At the Password: prompt, enter admin⁴. 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 89).



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.
 - 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 1999-2020

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 9.0.0.x at the 000000 build level.

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 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”](#) on page 92 to access the IBM MQ Appliance web UI and accept the license agreement.

M 2001 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 90 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](#).

M 2001 Maintenance

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

M 2001 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.

Related concepts

“Safety” on page 65

Before you install this product, read the Safety Information.

M 2001 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.

M 2001 LEDs on the front of the appliance

The following figure describes LEDs of the appliance.

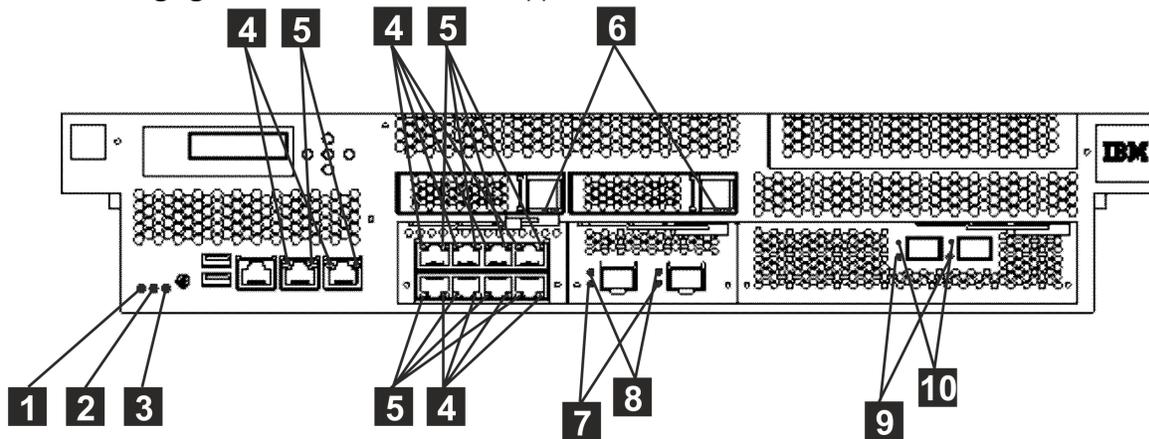


Figure 22. LEDs on the front of the M2000 appliance

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.

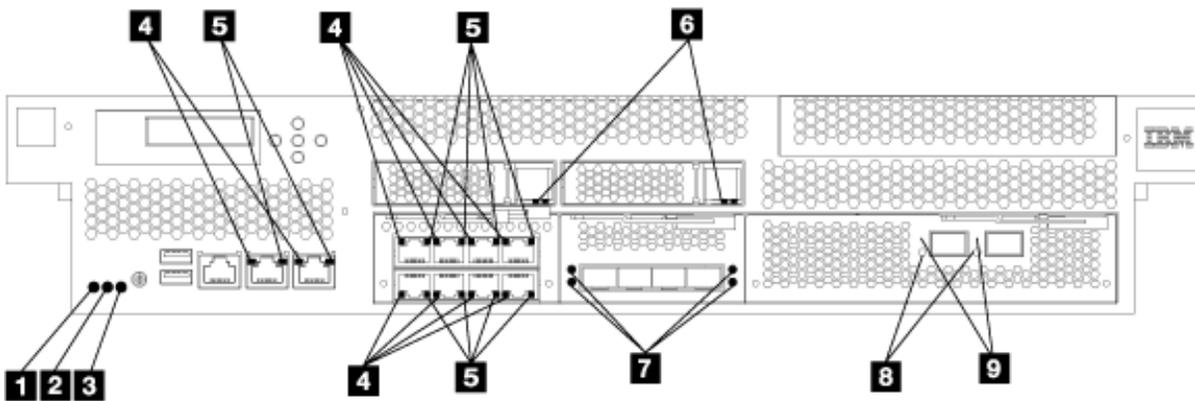


Figure 23. LEDs on the front of the M2001 appliance

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.

M 2001 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.

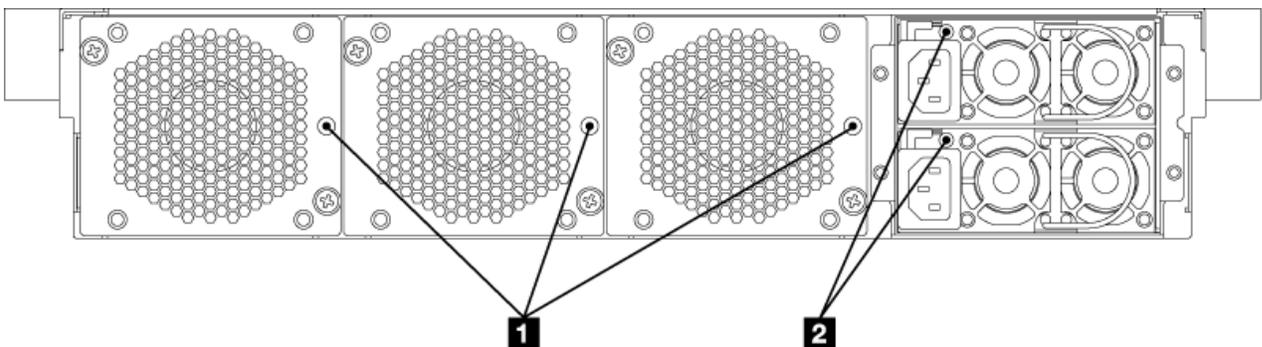


Figure 24. LEDs on the rear of the appliance

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.

M 2001 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.

M 2001 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, 2020 - 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**.

M 2001 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.

M 2001 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” on page 99](#).

No

Continue to step [“2” on page 99](#).

2. Does the log file contain a critical message?

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

Yes

Continue to step [“3” on page 99](#).

No

Continue to step [“4” on page 99](#).

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” on page 100](#).
- If the power supply module, see [“Troubleshooting power supply modules” on page 100](#)
- If the hard disk drive module, see [“Troubleshooting disk drive modules” on page 101](#).
- If field replaceable unit (FRU) parts, contact IBM Support.

No

Continue to step [“4” on page 99](#).

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

Yes

Continue with step [“5” on page 99](#).

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” on page 100](#).

If the power supply module, see [“Troubleshooting power supply modules” on page 100](#)

If the hard disk drive module, see [“Troubleshooting disk drive modules” on page 101](#).

No

The problem is with the appliance, use the appliance troubleshooting procedure.

M 2001 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.

M 2001 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.

M 2001 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.

M 2001 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.

Related concepts

[“test hardware command” on page 96](#)

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

Related tasks

[“Using the diagnostic self-test” on page 97](#)

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

M 2001 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*.

M 2001 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.

Related concepts

“[Guidelines for handling static-sensitive devices](#)” on page 102
Read these guidelines before you handle static-sensitive devices.

“[Safety statements](#)” on page 67
Safety statements are available on the included CD-ROM.

▶ M 2001 **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.

▶ M 2001 **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.

▶ M 2001 **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*.

▶ M 2001 **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.

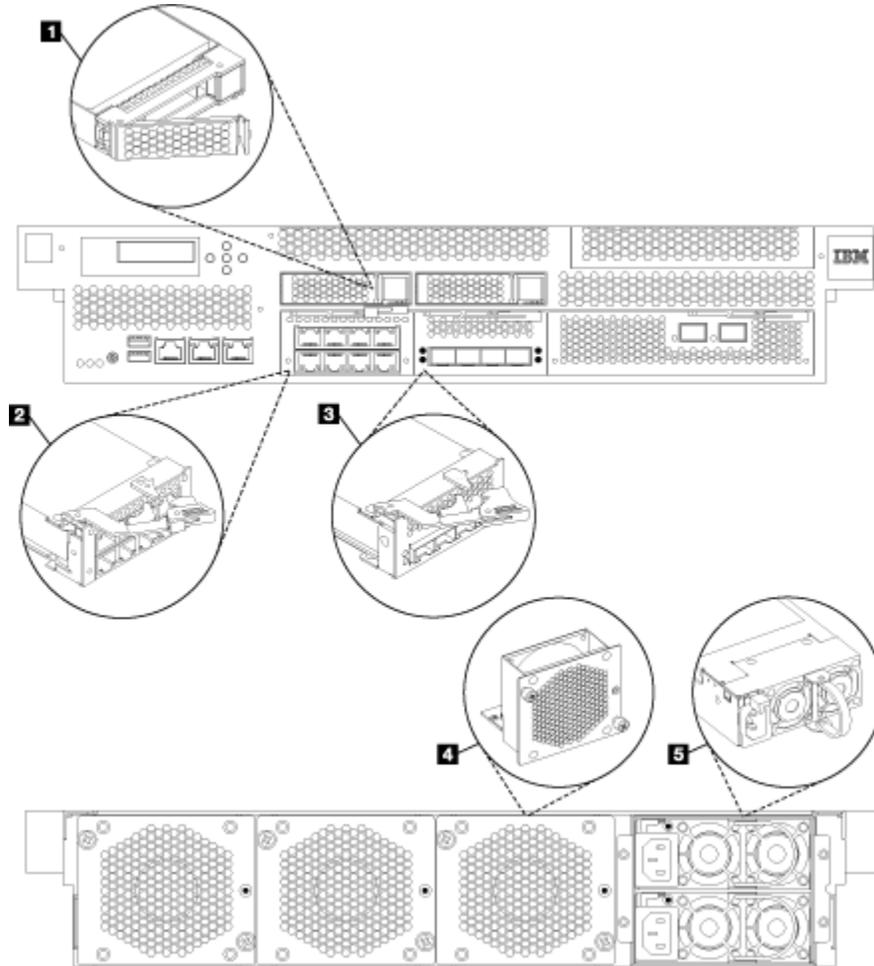


Figure 25. CRU part locations for the 8436 appliance.

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

Table 14. 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.		
-	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

Table 14. Part numbers for the IBM MQ Appliance. (continued)

Label	Description	Tier 2 CRU part number
-	SFP+ SR transceiver	46N5592

M 2001 *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 15. FRU part numbers for the appliance

Description	Part number
Shipping box	00VM076
Full MQ appliance system	01LK676
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

M 2001 **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 16. 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

Table 16. Power cords and cords (continued)

Country	Tier 2 CRU part number	Description
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
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

Related concepts

“Removing or replacing the appliance or parts” on page 101

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

M2001 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 poweroff** command to shut down the appliance.

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.

Related concepts

“LEDs on the front of the appliance” on page 93

The following figure describes LEDs of the appliance.

“Power button” on page 74

The front of the appliance has a power button.

M 2001 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” on page 106](#)
- [“Replacing a power supply module” on page 108](#)
- [“Replacing a solid state disk drive module - M2001 appliances” on page 110](#)
- [“Replacing an Ethernet module” on page 112](#)

M 2001 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.

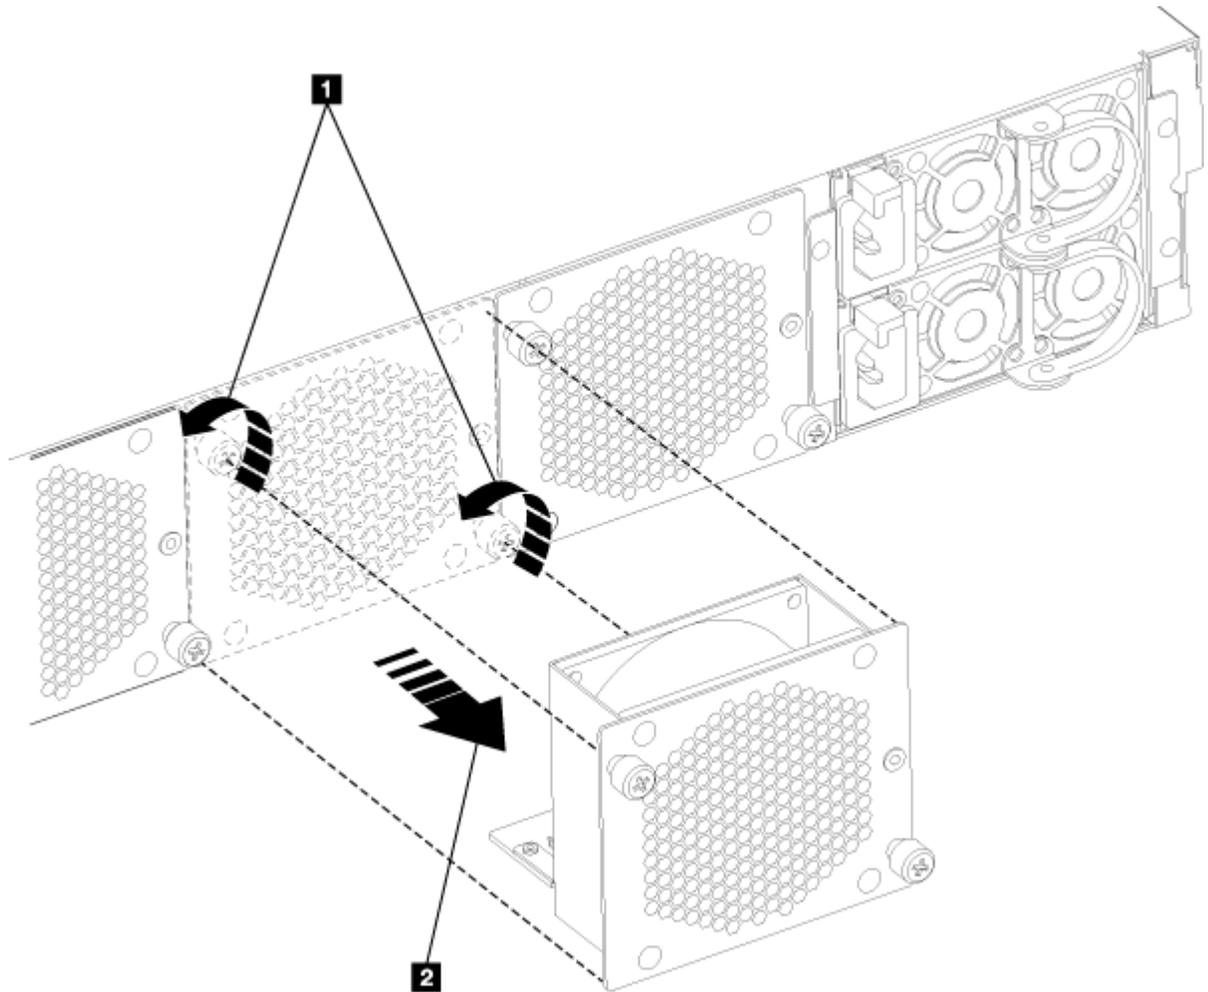


Figure 26. Removing a fan module

- 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 statements 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.

Related concepts

[“Fan modules” on page 80](#)

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

[“Returning an appliance or part” on page 102](#)

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

M 2001 *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.

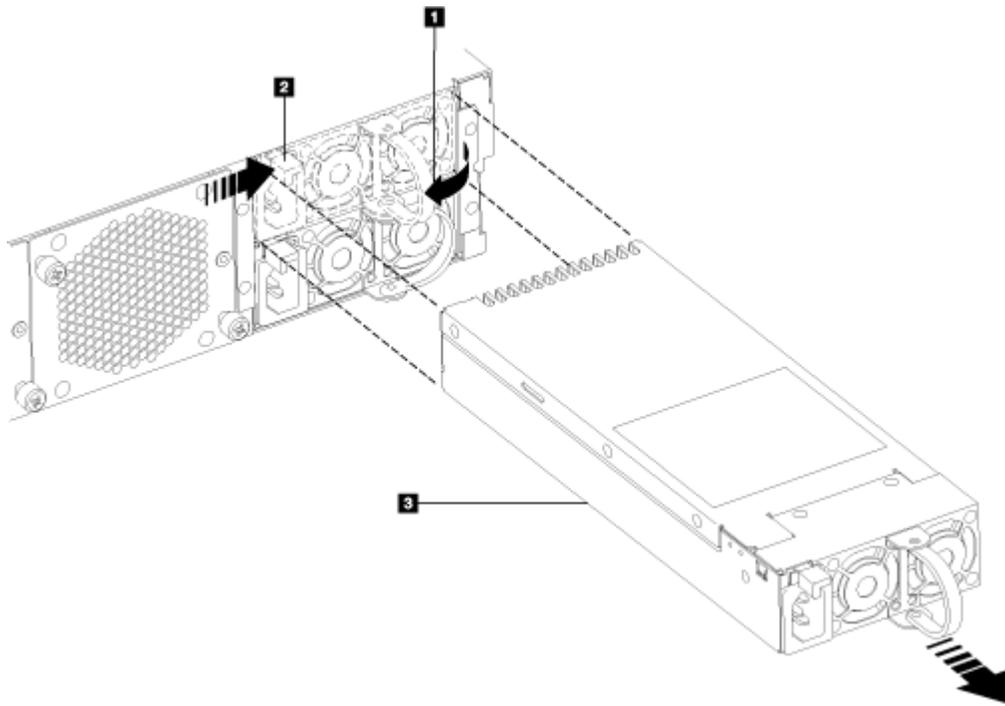


Figure 27. Removing a power supply module.

- 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.

Related concepts

[“Power supply modules” on page 80](#)

The appliance is powered by two redundant power supply modules.

[“Returning an appliance or part” on page 102](#)

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

M 2001 *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.

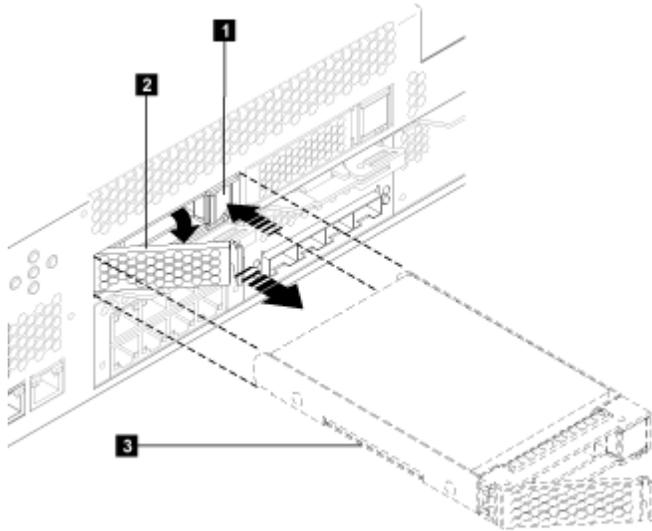


Figure 28. Removing a solid state disk drive module.

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.

Related concepts

[“Solid-state disk drive modules” on page 79](#)

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

[“Returning an appliance or part” on page 102](#)

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

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.

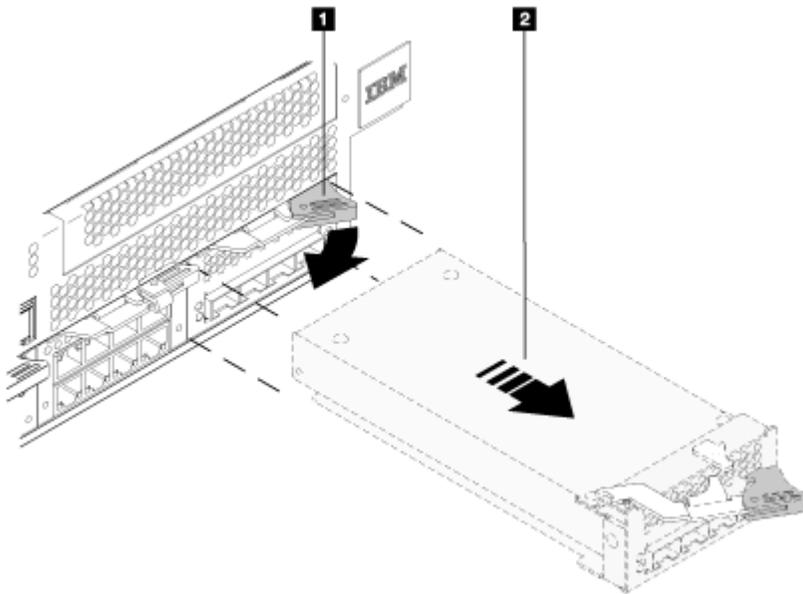


Figure 29. Removing the 10 Gb Ethernet module (M2001 model).

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.

Related concepts

[“Ethernet modules” on page 77](#)

The appliance contains two Ethernet modules for network connectivity.

[“Returning an appliance or part” on page 102](#)

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

Removing an SFP+ transceiver

The procedure to remove a 10 Gb SFP+ transceiver.

Before you begin

- The part number of a short reach transceiver module is 46N5592.
- The part number of a long reach transceiver module is 46N5593.

About this task



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. Wait until the power LED is no longer illuminated.
2. Unplug all power cords.

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

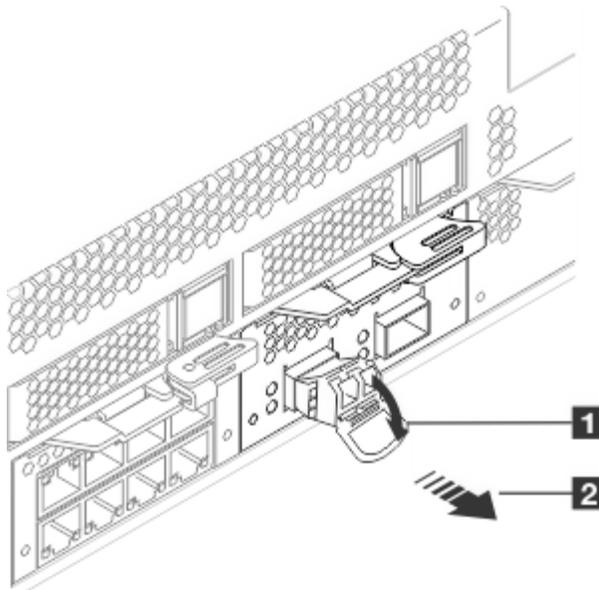


Figure 30. Removing the SFP transceiver

3. Pull downward on the latch at the front of the transceiver **1**.
4. Pull the transceiver out of the appliance by pulling forward on the release latch **2**.

M 2001 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.

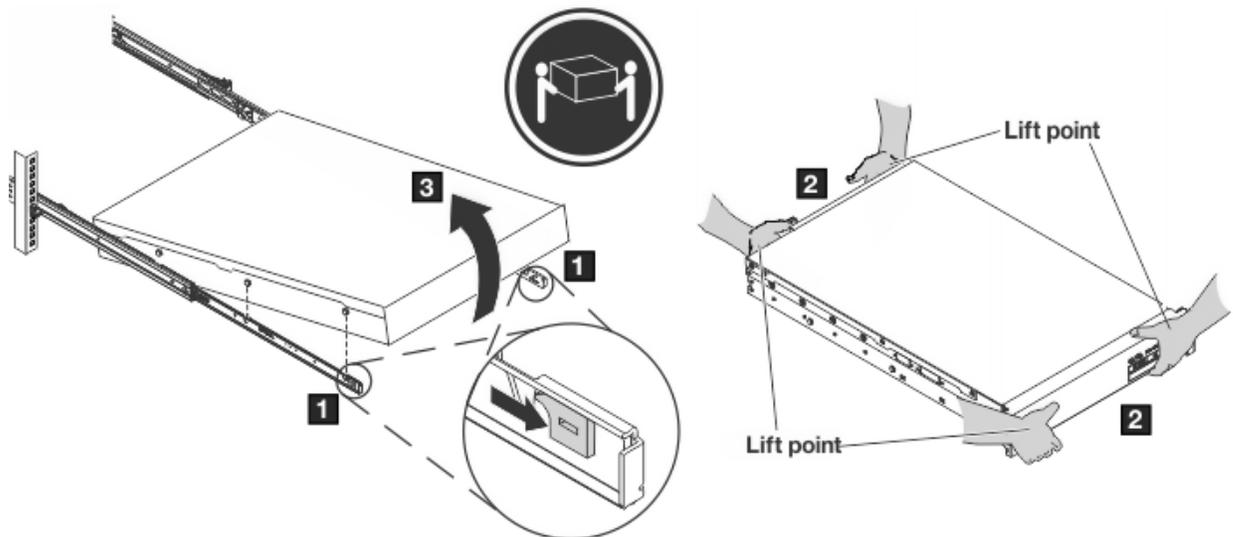


Figure 31. Unlatching and rotating the front of the appliance.

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.

Related concepts

[“Returning an appliance or part” on page 102](#)

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

M 2001 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)



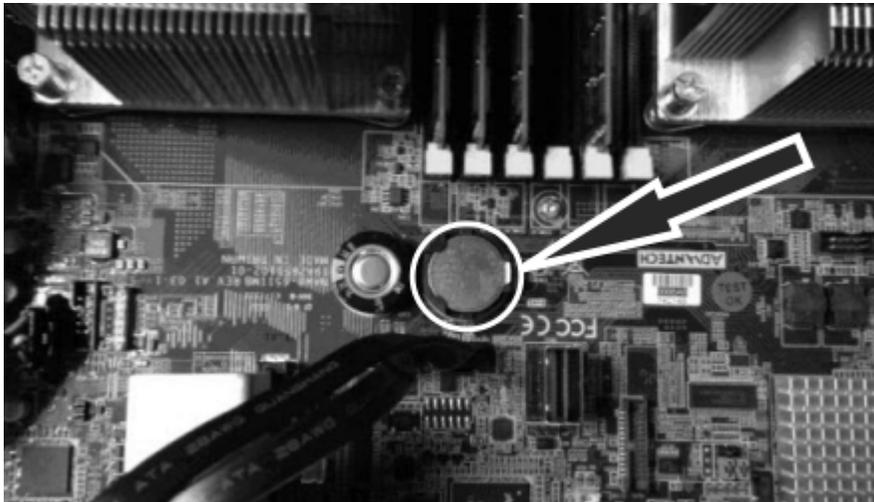
DANGER: Improper disposal or incineration of batteries or capacitors can cause life-threatening injury.

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.

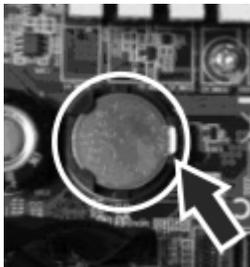
Procedure

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.

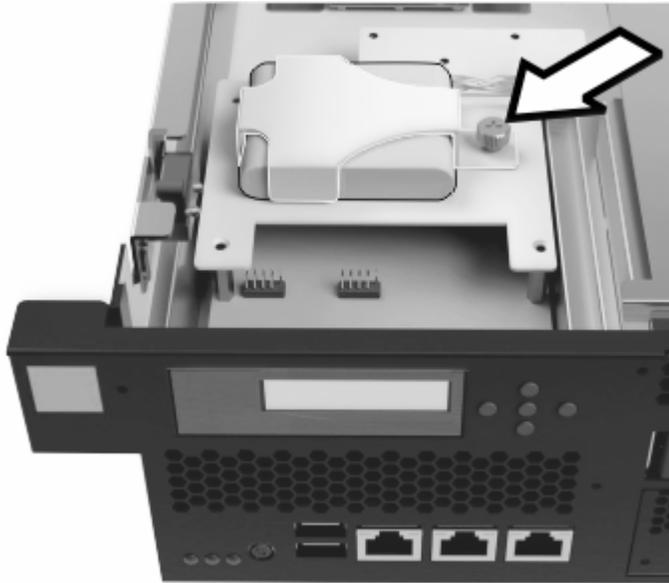
The battery is next to the RAM slots.



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.

M 2002 Installing the M2002 appliance

Plan for your installation, install, and verify the installation of the IBM MQ Appliance, model M2002.

Safety

Before you install this product, read the Safety Information.

Arabic

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الأمنية.

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 ασφαλείας (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.

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 pečítajte 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

Safety notices apply to this product. The safety notices include danger and caution notices.

Danger notice definition

A special note in the text that calls attention to a situation that is potentially lethal or extremely hazardous to people.

Caution notice definition

A special note in that text that calls attention to a situation that is potentially hazardous to people because of some existing condition, or to a potentially dangerous situation that might develop because of some unsafe practice.

Danger notices

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 have completed the voltage and grounding checks. Improper wiring or grounding could place dangerous voltage on the metal shell. If any of the conditions are not as described, *STOP*. Ensure the proper voltage or impedance conditions are corrected before proceeding. (D003)



DANGER: An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or devices that attach to the system. It is the responsibility of the customer 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:

- If IBM supplied power cords, 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.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening 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 the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
2. Attach all cables to the 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

Caution notices apply to this product.



CAUTION: The battery contains lithium. To avoid possible explosion, do not burn or charge the battery. Do not throw or immerse into water, heat to more than 100° C (212° F), repair or disassemble. (C003).



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)



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 those specified herein might result in hazardous radiation exposure. (C026)



CAUTION: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. (C027)



CAUTION:

- Do not install a unit in a rack where the internal rack ambient temperature will exceed the manufacturer's recommended ambient temperature for all 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 unit used for air flow through the unit.
- Consideration should be given 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 located on the equipment in the rack to 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.



- (For fixed drawers): This drawer is a fixed drawer and 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 of 2)

Labels

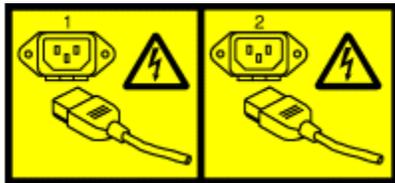
One or more of the following safety labels may apply to this product.



 **DANGER:** Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)



 **DANGER:** Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). (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 M2002

The IBM MQ Appliance provides IBM MQ V9.1 on an easy-to-deploy appliance.

The IBM MQ Appliance has the model number M2002, and the MTM 8441-54X.

Specifications and features

Information about the specifications and hardware features of appliances.

Hardware specifications

The following table summarizes the specifications for the chassis.

<i>Table 17. Appliance specifications</i>	
Dimensions	
Height	3.46 in. (88 mm)
Width	18.98 in. (482 mm)

Table 17. Appliance specifications (continued)

Depth	26.97 in. (685 mm)
Appliance weight	69.5 lb. (31.5 kg)
Shipping weight	91.5 lb. (41.5 kg)
Electrical input	
Power Supply	Two, 1200 Watt power supply modules
PSU Receptacle Type	PSUs have C13 inlets (two redundant PSUs)
Sine-wave	47/63 Hz (single-phase) required
110 Voltage AC	100 to 127 Volt (nominal) at 12.47 Amperes
220 Voltage AC	200 to 240 Volt (nominal) at 7.08 Amperes
Max PSU Rating (Watts)	Redundant 1200W PSUs
Maximum Power Consumption (Watts)	850W
Typical Power Consumption (Watts)	~650W
Heat output	
Idle	222.88 watts (761.25 Btu) per hour
Maximum	522.23 watts (1783.7 Btu) per hour
Airflow direction	Front to back
Environment	
Shipping	-40° to 140° F (-40° to 60° C)
Power off	-4° to 140° F (-20° to 60° C)
Power on	50° to 95° F (10° to 25° C)
Maximum altitude	Operating: 6998 ft. (2133m) Storage: 39,370 ft. (12,000m)
Humidity	40° C (104° F) at 95% relative humidity (operating)

Acoustic

The following table shows the acoustic specifications.

Table 18. Noise emissions

Product description	Declared A-weighted sound power level, L _{WA} d (B)		Declared A-weighted sound pressure level, L _{pAm} (dB)	
	Operating	Idling	Operating	Idling
M2002	7.8	Not tested	62	Not tested

Table 18. Noise emissions (continued)

Product description	Declared A-weighted sound power level, L_{WA_d} (B)		Declared A-weighted sound pressure level, L_{pA_m} (dB)	
	Operating	Idling	Operating	Idling
<p>Notes:</p> <ol style="list-style-type: none"> 1. Declared level L_{WA_d} is the upper-limit A-weighted sound power level. Declared level L_{pA_m} is the mean A-weighted sound pressure level measured at the 1-meter bystander positions. 2. All measurements made in conformance with ISO 7779 and declared in conformance with ISO 9296. 3. B, dB, abbreviations for bels and decibels, respectively. 1 B = 10 dB. 4. <p>Notice: Government regulations (such as those prescribed by OSHA or European Community Directives) may govern noise level exposure in the workplace and may apply to you and your server installation. This IBM system is available with an optional acoustical door feature that can help reduce the noise emitted from this system. The actual sound pressure levels in your installation depend upon a variety of factors, including the number of racks in the installation; the size, materials, and configuration of the room where you designate the racks to be installed; the noise levels from other equipment; the room ambient temperature, and employees' location in relation to the equipment. Further, compliance with such government regulations also depends upon a variety of additional factors, including the duration of employees' exposure and whether employees wear hearing protection. IBM recommends that you consult with qualified experts in this field to determine whether you are in compliance with the applicable regulations.</p>				

Hardware features

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

Table 19. Appliance features		
CPU	Disk space	Memory
Two 12-core 2.60 GHz Intel Xeon Gold 6126T processors	Four 3200 GB SSDs configured as RAID10	192 GB (Twelve 2666 MHz DDR4 DIMMs)

The system disk contains 32 GB space for system file storage.

The RAID array for user storage on model M2002A contains 6 TB of storage, and 3 TB for model M2002B. Allocation of storage is set during appliance initialization. The RAID array is formatted as a single partition.

► M2002 Intrusion detection

The appliance has an internal intrusion detection switch. 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 to reset intrusion detection.

If intrusion detection is enabled and the appliance detects an intrusion during normal operation, a warning message is displayed in the GUI for newly connected sessions and the appliance restarts in failsafe mode.

An administrator can reset the intrusion detection with the **clear intrusion-detected** command.

M 2002 Identifying components

To work with the product, you must understand the components on the front and rear on the appliance.

Components on the front

Controls, connectors, and status indicators are at the front of the appliance. The following figure shows the components at the front of the appliance.

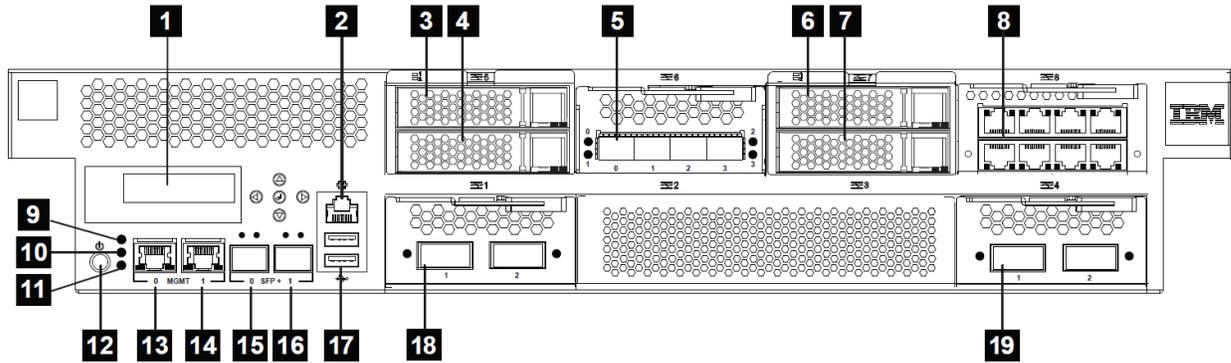


Figure 32. Front view

- 1** LCD display.
- 2** Console connector.
- 3** Solid-state disk drive 1.
- 4** Solid-state disk drive 2.
- 5** 10 Gb Ethernet module.
- 6** Solid-state disk drive 3.
- 7** Solid-state disk drive 4.
- 8** 1 Gb Ethernet module.
- 9** Fault LED.
- 10** Locate LED.
- 11** Power LED.
- 12** Power button.
- 13** 1 Gb Management Ethernet port 0.
- 14** 1 Gb Management Ethernet port 1.

- 15** 10 Gb Ethernet port 0.
- 16** 10 Gb Ethernet port 1.
- 17** Two USB ports.
- 18** 40 Gb Ethernet module.
- 19** 40 Gb Ethernet module.

Components on the rear

Fans and power supplies are on the rear of the appliance. The following figure shows the components on the rear of the appliance.

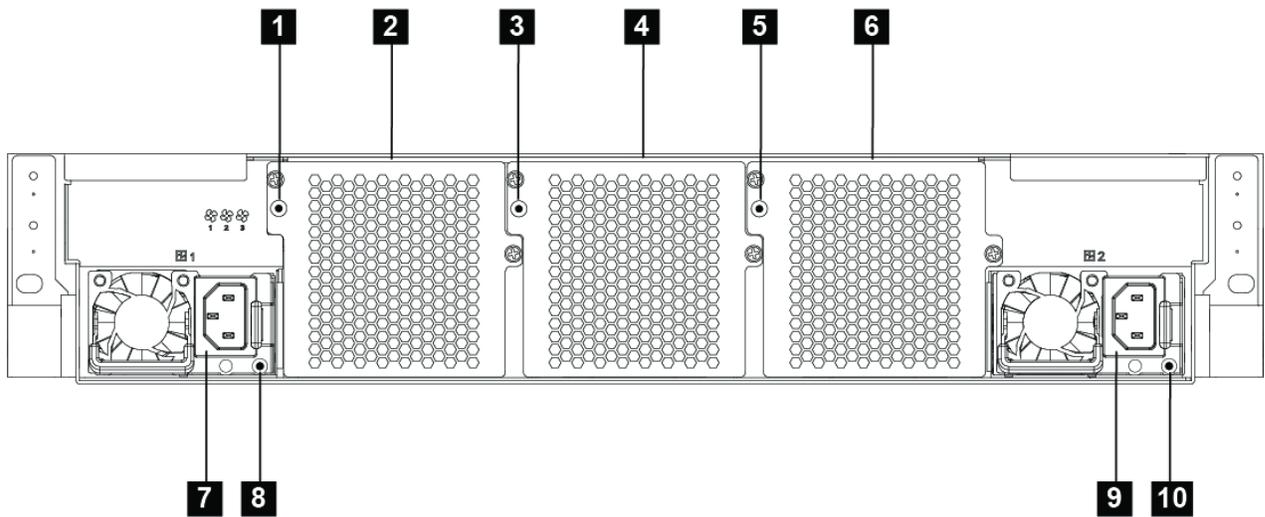


Figure 33. Rear view

- 1** LED for fan module 1
- 2** Fan module 1
- 3** LED for fan module 2
- 4** Fan module 2
- 5** LED for fan module 3
- 6** Fan module 3
- 7** Power supply module 1
- 8** LED for power supply module 1
- 9** Power supply module 2

10

LED for power supply module 2

M 2002 LCD module

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

The LCD provides information about the appliance. The menu buttons are not functional.

M 2002 Locate LED

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

The locate LED is steady blue when activated and remains on until deactivated.

From the GUI

1. In the search field, enter **system**.
2. From the results, click **System Control**.
3. Locate the **Control Locate LED** section.
 - To activate, click **on**.
 - To deactivate, click **off**.
4. Click **Control Locate LED**.

From the CLI

Use the Global **locate-device** command.

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

M 2002 Power button

The front of the appliance has a power button.

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

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

Related tasks

[“Turning off the appliance” on page 157](#)

How to turn off power to the appliance.

M 2002 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 a PC that is running terminal emulation software to the appliance or to connect from an ASCII terminal to the appliance. An ASCII terminal is a simple device that transmits and receives ASCII data.

USB ports

The appliance has two USB ports that conform to USB 3.0 devices.

The USB ports are not active. Therefore, they do not provide any connection.

Network ports

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

Network ports are grouped and located by function.

- Two fixed management Ethernet ports.
- Two fixed 10 Gb small-form-factor pluggable SFP+ ports.
- One removable 1 Gb Ethernet module with eight ports for the RJ45 interface.

- One removable 10 Gb Ethernet module with four SFP+ ports.
- Two removable 40 Gb Ethernet modules, each with two QSFP+ ports.

Management Ethernet ports

The management Ethernet ports provide remote management access to the appliance and should not be used as data ports. Use the ports in the Ethernet modules for handling data traffic.

Consider management traffic 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.

mgt0

This Ethernet interface can manage all management data. This port also supports IPMI over LAN, which includes serial over LAN.

mgt1

This Ethernet interface can manage all management data. This port also supports IPMI over LAN, which includes serial over LAN.

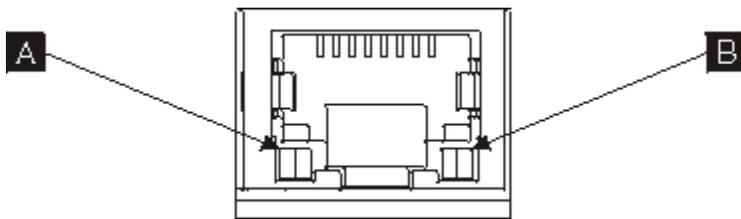


Figure 34. Management port LEDs

A

Speed LED

B

Activity LED

10 Gb Ethernet ports

The appliance has two 10 Gb small-form-factor pluggable SFP+ ports that are part of the appliance. The SFP+ ports support optical or electrical interfaces with the appropriate transceiver. These two ports are functionally the same with the ports in the 10 Gb Ethernet module.

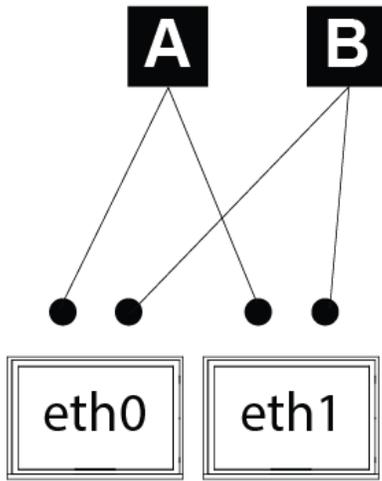


Figure 35. 10 Gb Ethernet port LEDs

A

Speed LED

B

Activity LED

Ethernet modules

The appliance has three Ethernet modules for network connectivity. The right module contains eight 1 Gb Ethernet ports. The middle modules contain four 40 Gb Ethernet ports. The left module contains four 10 Gb Ethernet ports.

1 Gb Ethernet module

Has eight RJ45 Ethernet ports. The Ethernet ports are placed in two rows. The numbering start with 10 and are sequential from lower-left to upper-right. Each port has an associated speed LED and activity LED. The speed and activity LEDs on the lower and upper rows have opposite orientation.

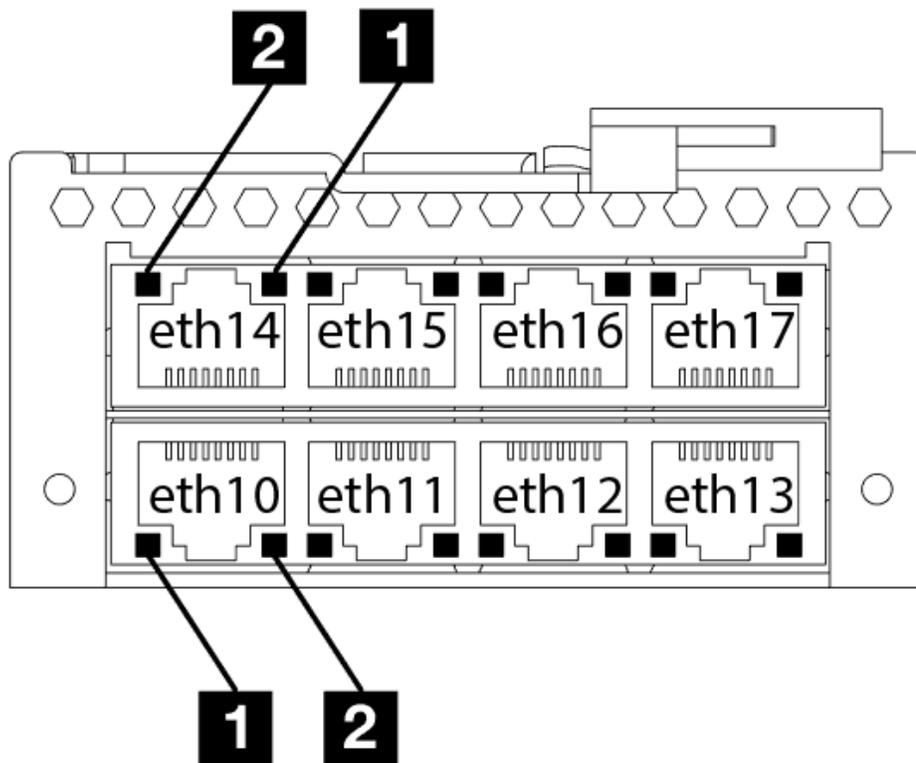


Figure 36. 1 Gb Ethernet module LEDs

- 1** Activity LED
- 2** Speed LED

10 Gb Ethernet module

Has four 10 Gb SFP+ ports. The numbering start with 20 and are sequential from left to right. Each port has an associated activity LED. By default, eth20 is used for the replication link in a disaster recovery configuration and eth21 is used for the replication link in a high availability configuration.

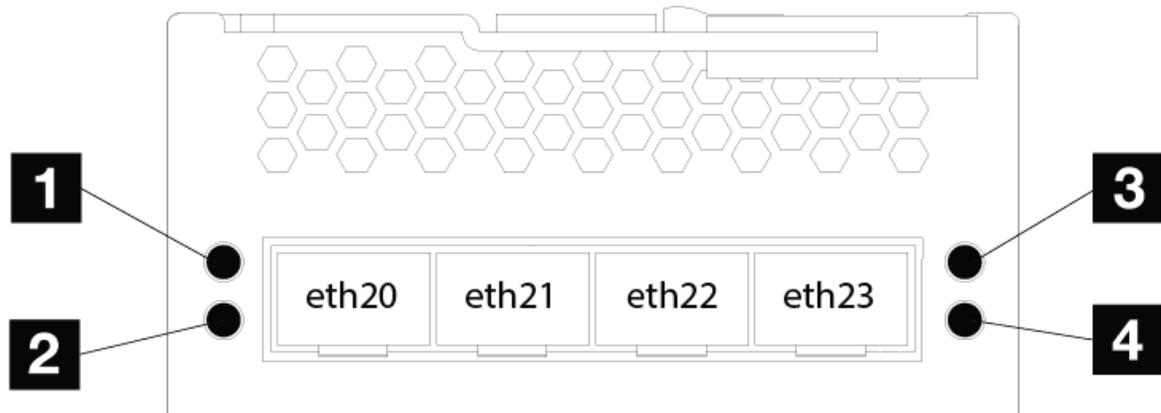


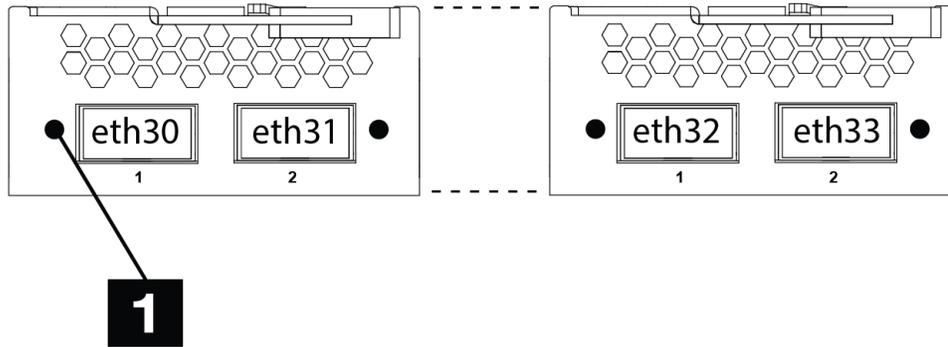
Figure 37. 10 Gb Ethernet module LEDs

- 1** Activity LED for eth20

- 2** Activity LED for eth21
- 3** Activity LED for eth22
- 4** Activity LED for eth23

40 Gb Ethernet modules

Has two 2 x 40 Gb QSFP+ ports. They are numbered 30, 31, 32, 33 respectively. Each port has an associated activity LED.



- 1** Activity LED for eth30

Each 40 Gb Network module provides a total of 40Gb of bandwidth that is shared between the two physical ports on that module. You can create an aggregate interface using one physical connection to each network module (for example, eth30 and eth32) to boost total bandwidth and the availability characteristics of the link.

Solid-state disk drive modules

The M2002 appliance has four solid-state disk drive modules.

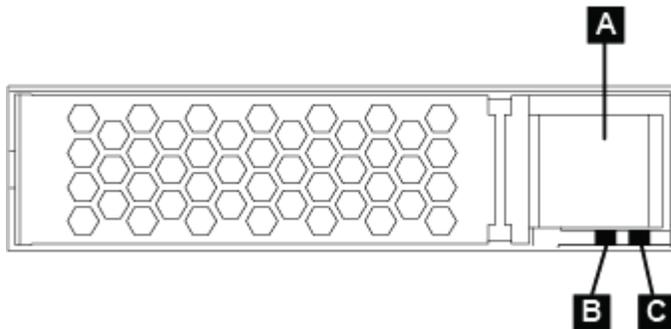


Figure 38. Solid-state disk drive module.

- A** Locking arm release latch.
- B** Activity LED.
- C** Fault LED.

Fan modules

The appliance has three fan modules. Each fan module contains a cooling fan with an LED that indicates the status of the module.

- Single flash when power is first applied.

- Steady when the fan is operating at less than 1200 RPM or when there is a fault.
- No illumination when operating normally or when no power is present.

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

Power supply modules

The appliance is powered by two redundant power supply modules.



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

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.

- Steady green indicates that the module is connected to a power source.
- Steady red indicates that the module is not functioning within design specifications.
- Not illuminated indicates that there is no power to the module.

Related tasks

[“Turning off the appliance” on page 157](#)

How to turn off power to the appliance.

▶ M 2002 Installation preparation

Information about the rack and tools that you need to prepare for installation.

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 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).

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 as many network cables as there are network ports to connect the appliance to your network.

▶ M 2002 Installation of the appliance in a rack

The information that you need before you install the appliance in a rack.

The rails are for a 19 in. (48.26 cm) rack. A complete rail kit is required to install the appliance.

The rail kit contains the following parts:

- Left slide rail, marked L.
- Right slide rail, marked R.

- Two (2) size 10-32 screws to secure the slide rails to the rack.

If any item is missing, contact IBM Support.

M 2002 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.

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.

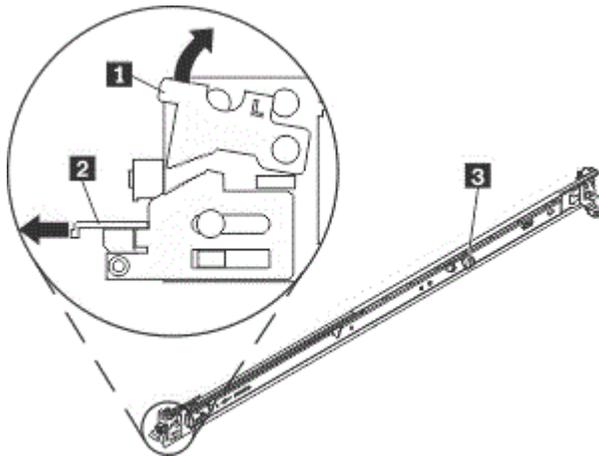


Figure 39. View of the left slide rail.

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.

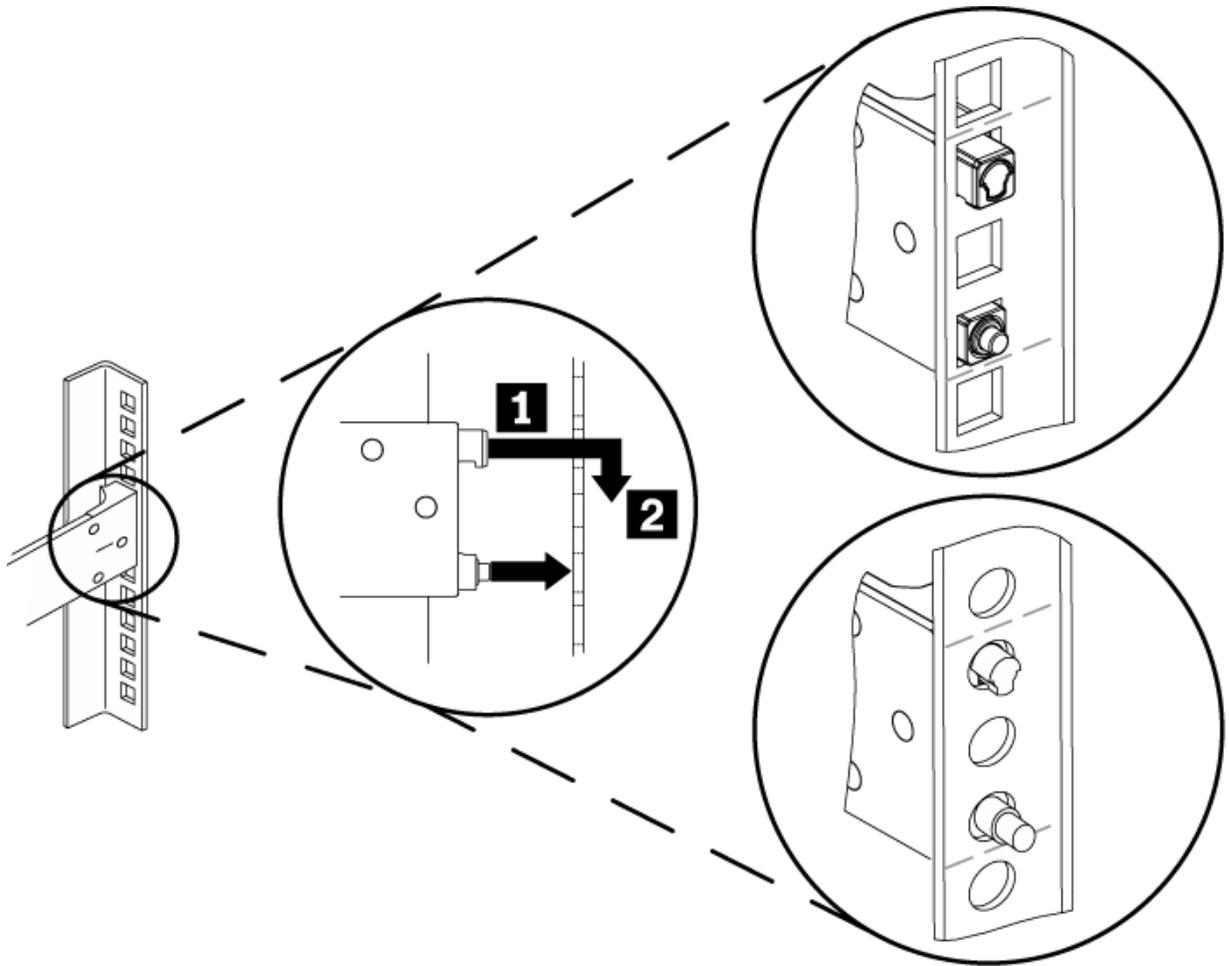


Figure 40. Install the rear end of the slide rails.

- 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.

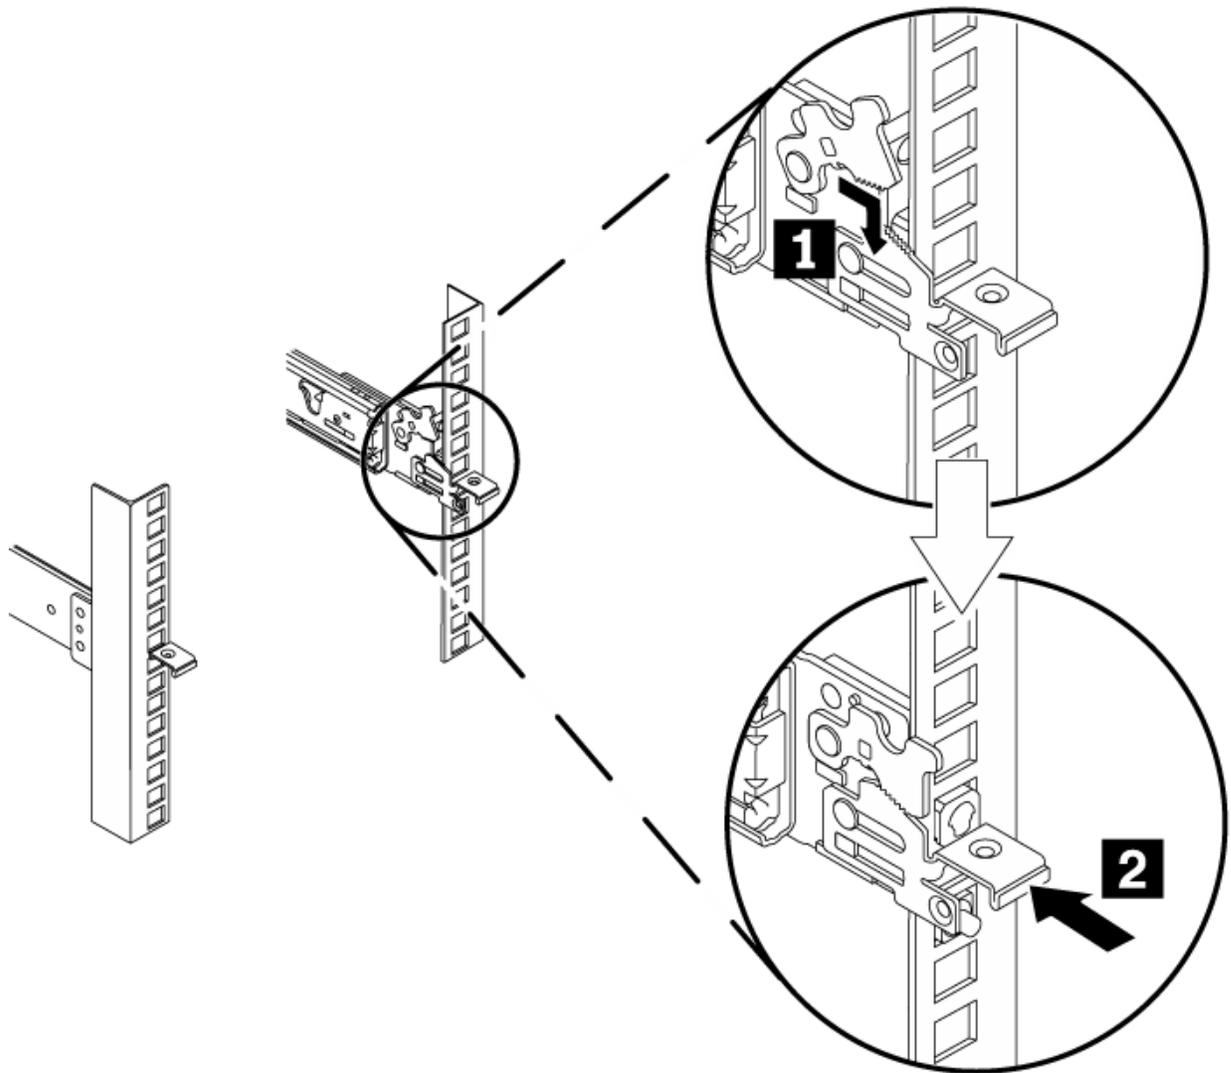


Figure 41. Install the front end of the slide rails.

- 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.

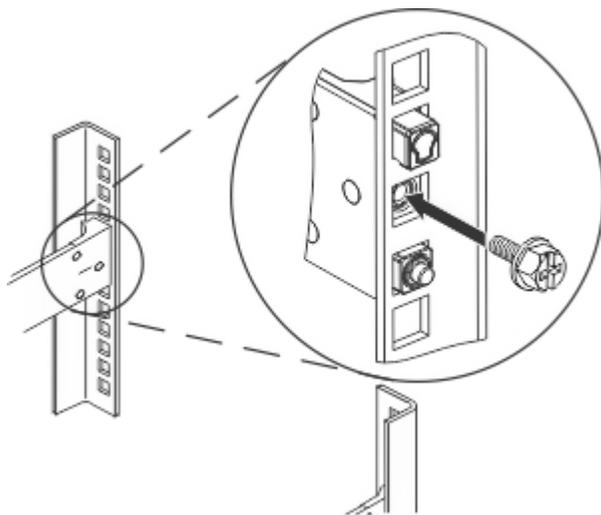


Figure 42. Securing the rails in the rack.

6. Repeat step 5 for the left rail.

M 2002 Installing the appliance on the rails

How to install the appliance on the rails.

About this task



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. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). (L002)

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

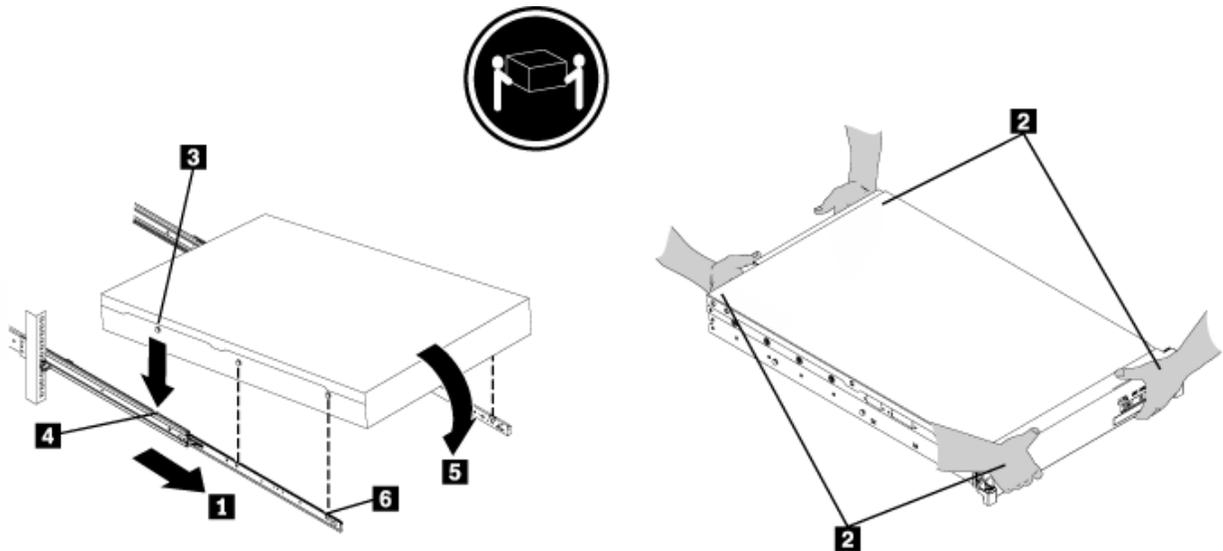


Figure 43. Securing the appliance in the rack

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.

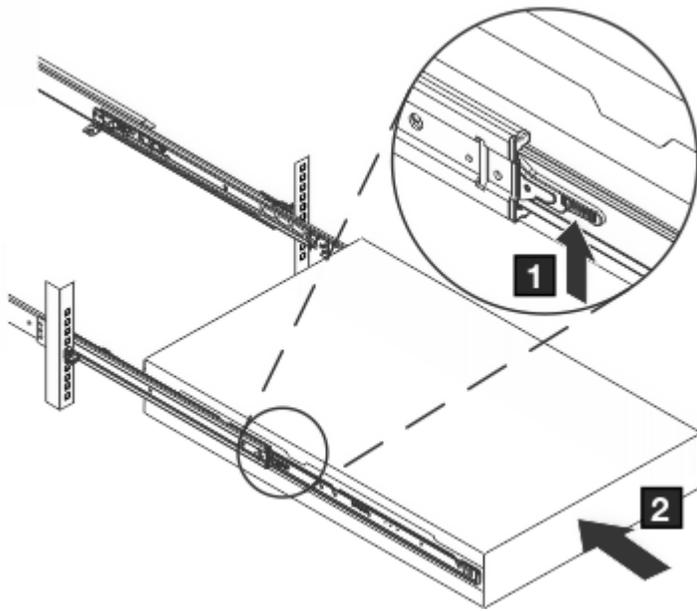


Figure 44. Sliding the appliance into the rack.

Procedure

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

M 2002 Considerations when connecting the appliance to AC power sources

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 have completed the voltage and grounding checks. Improper wiring or grounding could place dangerous voltage on the metal shell. If any of the conditions are not as described, *STOP*. Ensure the proper voltage or impedance conditions are corrected before proceeding. (D003)



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

Use the provided power cords to connect both power supply modules to AC power sources. You must connect each power supply modules, or the system considers the unconnected module to be in a failed state.

Each power supply can come from different power sources or from different circuits. The two sources must be properly grounded.

M 2002 Considerations when connecting the appliance to your network

Considerations to review before you connect the appliance to your 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: 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 those specified herein might result in hazardous radiation exposure. (C026)



CAUTION: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, 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.

The 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.

10Gb and 40Gb network connectivity is provided through network modules that require an external transceiver (10Gb SFP+, 40Gb QSFP+) appropriate to your chosen physical network. IBM provides short range (SR) optical transceivers with the appliance. These transceivers conform to the following specifications.

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

40GBASE (40 Gbps) connection:

Short-reach (300 meters) QSFP+ modules with MPO connector (multi-mode fiber)

- Optical interface specifications per IEEE 802.3ba 40GBASE-SR4

All optical transceivers provided by IBM are Class 1 Eye safe per requirements of IEC 60825-1 /CDRH.

A DAC (Direct attach copper) 10GB cable is also supplied with each appliance. This cable can be used in place of the supplied optical transceivers, for example, if directly connecting two appliances for high availability replication.

You might choose to replace the supplied transceivers in the 10GB or 40GB network ports with alternative parts, for example, to replace a short range optical transceiver with a long range optical transceiver. Any SFP+ and QSFP+ transceiver (or DAC/AOC) following industry standards is compatible with the IBM MQ Appliance network modules. If a support problem is identified with an SFP+ or QSFP+ transceiver not supplied by IBM, you might be asked to reproduce the problem on a component supplied by IBM. The 'link' and/or 'activity' LEDs can be used to confirm that the transceiver is operational.

The network modules in the IBM MQ Appliance are based on the following Intel adapters, and any transceivers known to be compatible with these adapters are therefore appropriate:

- 4-port SFP+ 10GbE NIM is equivalent to the Intel Ethernet Converged Network Adapter X710-DA4
- 2-port QSFP+ 40GbE NIM is equivalent to the Intel Ethernet Converged Network Adapter XL710-QDA2

▶ M 2002 **Setting up the initial firmware configuration**

How to perform the base, initial firmware configuration.

About this task

This configuration is the minimal configuration to add an appliance to your environment. Defining the full configuration is beyond the scope of this document.

Procedure

1. Read the hardware and information requirements, and read the considerations for the operation modes and the password for the `admin` account.
2. Connect the serial cable.
3. Initialize the appliance by changing the password for the `admin` account and interactively define the initial configuration.
4. Accept the license agreement and verify the configuration.

▶ M 2002 **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.

- Review and comply with the hardware requirements.

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.

- Obtain the required network data.
 - The IP addresses for the Ethernet interfaces for appliance management.
 - The IP addresses for the Ethernet interfaces for service access.
 - The IP addresses of default gateways (routers) that support the subnets for Ethernet interfaces.
 - The IP addresses and ports for the web management interface and SSH service.
 - Optional: The IP address and port for Telnet service.

Tip:

- The web management interface is required to accept the license agreement.
- To use an IPMI connection, which includes serial over LAN, requires the configuration of the `mgmt0` interface.

▶ M 2002 **Installation considerations**

During the initial configuration, you are prompted for supported operational modes and the password for the `admin` account.

Operational modes

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

- IBM MQ Appliance M2002A is aimed at larger enterprise workloads.
- IBM MQ Appliance M2002B is designed to meet the needs for smaller workloads and offers a lower processing capability and less disk capacity.

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 mode 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 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 M2002B appliance to an M2002B+ appliance, which has the same capacity as an M2002A appliance.

Upgrading from an M2002B to M2002B+ 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.

V 9.1.4 You can view which mode is currently operational on the appliance in the following ways:

- By viewing the LCD display panel on the front of the appliance.
- In the welcome banner when you start the command line interface.
- By using the **show system** command (reported as Capacity mode). See [“show” on page 885](#).
- In the System Setting panel of the appliance Web UI (displayed as **Licensed capacity mode**).
- By using a REST API command to show system settings. See [“REST management interface” on page 1000](#).

Password for the admin account

The first time that you start the appliance from a serial connection is different from any subsequent restart.

- On the first start, you must initialize the appliance. The initialization routine prompts you to accept the license agreement and change the password for the admin account. After initialization, create a user with a group-defined account type with the appropriate access policy or the privileged account type as a backup for the admin account. A privileged user or a group-defined user with the appropriate access policy can log in and reset the password for the admin account.
- On subsequent restarts, you are prompted for the credentials of the admin account or another local account. If that account has an expired password, you are prompted to change the password.



Attention: If you forget or misplace the password for the admin account, you must return the appliance to IBM to reset the password. However, when another account can log in and has appropriate permission, that user can reset the password for the admin account.

When you receive the appliance after a password-reset, you must initialize the appliance again with an initial firmware setup. Therefore, none of your configuration data is on the appliance.

License agreement

During initialization, you must configure the web management service to prepare for accepting the license agreement through the GUI. When the admin account logs in to the GUI for the first time to verify the appliance configuration, the GUI displays the license agreement.

M 2002 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 142](#).

Read the operation modes and password considerations for the admin account in [“Installation considerations” on page 142](#).

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 Microsoft 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, 8N1⁵, and no flow control data transfer.

What to do next

Define the base configuration, which includes changing the password for the `admin` account, creating privileged users, and configuring the web management interface. See [“2 of 3: Initializing the appliance” on page 144](#).

M 2002 2 of 3: Initializing the appliance

How to initialize the appliance to define its initial configuration.

Before you begin

Connect the appliance to an ASCII terminal or to a PC that is running terminal emulation software through a serial connection. See [“1 of 3: Connecting the serial cable to the appliance” on page 143](#).

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` as the account name.⁶
3. At the `Password:` prompt, enter the password for the `admin` account.⁷ The script prompts you later to change this password.

⁵ 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.

⁶ `admin` is the name of a local user account. The owner of this account can perform all tasks on the appliance.

⁷ `admin` is the default password for the `admin` account except appliances that are delivered to Argentina. In Argentina, the password is provided on a separate CD.

4. Follow the prompts to enable the appropriate appliance mode. Select M2002A or M2002B according to your license, see [“Operational modes”](#) on page 142.



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.

If you inadvertently enter `n`, you can start the installation wizard by entering the following commands.

```
configure terminal
startup
```

8. Follow the prompts to complete the base firmware configuration.
 - Initializing the RAID array for the appliance can take a while.
 - To prepare for accepting the license agreement, you must configure the web management interface when prompted, or from the CLI with the `web-mgmt` command.

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, for example:

```
Welcome to IBM MQ appliance M2002A console configuration.
Copyright IBM Corporation 1999-2020
Version: MQ00.9.1.0.0 build 000000 on 2018/08/18 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#
```

What to do next

Access the GUI to accept the license agreement. See [“3 of 3: Accepting the license agreement”](#) on page 145.

3 of 3: Accepting the license agreement

You must access the GUI to accept the license agreement.

Before you begin

Define the base configuration for the appliance. See [“2 of 3: Initializing the appliance”](#) on page 144.

About this task

This procedure makes the following assumptions.

- The IP address for the Ethernet interface that is used to access the GUI is `10.10.13.35`
- The specialized HTTP server to support GUI access listens on port `9090`

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 GUI 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.

► M 2002 **Completing the configuration**

Information on completing the configuration is given in the Configuration section of the Knowledge Center.

[“Configuring” on page 207](#)

► M 2002 **Diagnostics and the appliance**

The appliance and its firmware provide various ways to help you diagnose problem with hardware components.

► M 2002 **M2002 LED diagnostics**

How to use LEDs on the M2002 appliance to diagnose problems.

LEDs can help you diagnose possible problems with hardware components.

LEDs on the front panel

The following figure describes LEDs on the front panel.

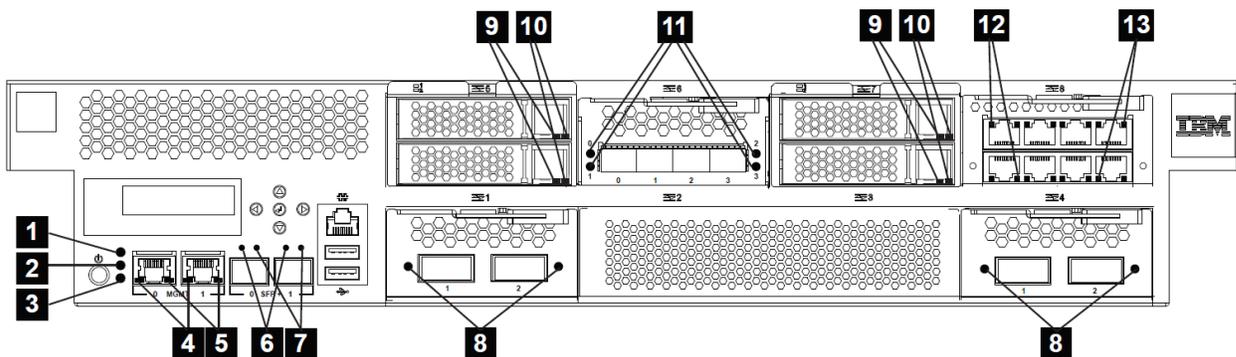


Figure 45. LEDs on the front of the appliance

The labels in this figure correspond to the following LEDs.

- 1** Fault LED
Steady amber when a critical hardware event is detected.
- 2** Locate LED
Steady blue when activated.

3

Power LED

Steady green indicates that power is connected and the appliance is turned on.

4

1 Gb Ethernet port activity LED

- Steady green indicates that the port is connected.
- Flashing green indicates port activity.

5

1 Gb Ethernet port speed LED

- Steady green indicates a 1 Gb Ethernet connection.
- Steady amber indicates a 10 or 100 Mbps connection.

6

10 Gb Ethernet port speed LED

Steady green indicates a 10 Gb Ethernet connection.

7

10 Gb Ethernet port activity LED

- Steady green indicates that the port is connected.
- Flashing green indicates that the transaction is ongoing.

8

40 Gb Ethernet port activity LED

- Steady green indicates that the port is connected.
- Flashing green indicates that the transaction is ongoing.

9

Solid-state disk drive activity LED

- Steady green indicates that the module is inserted fully.
- Flashing green indicates the reading or writing of data on the disk.

10

Solid-state disk drive fault LED

Steady amber indicates that the hard disk drive failed. This LED is nonfunctional.

11

10 Gb Ethernet port activity LED

- Steady green indicates that the port is connected.
- Flashing green indicates that the transaction is ongoing.

12

1 Gb Ethernet port speed LED

- Steady green indicates a 1 Gb Ethernet connection.
- Steady amber indicates a 10 or 100 Mbps connection.

13

1 Gb Ethernet port activity LED

- Steady green indicates that the port is connected.
- Flashing green indicates port activity.

LEDs on the rear panel

The following figure describes LEDs on the rear panel.

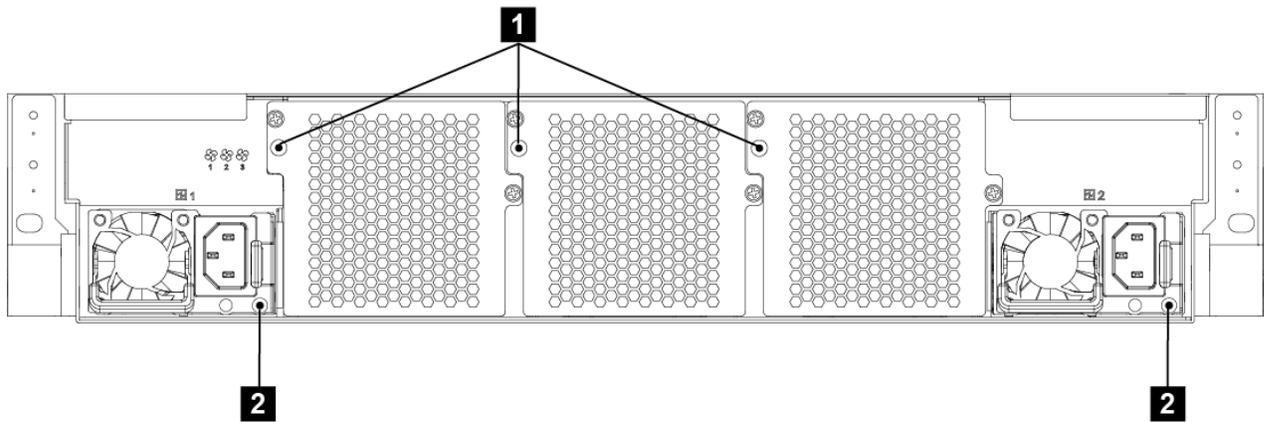


Figure 46. LEDs on the rear of the appliance

The labels in this figure correspond to the following LEDs.

1

Fan LEDs.

- Single flash when power is first applied.
- Steady when the fan is operating at less than 1200 RPM or when there is a fault.
- No illumination when operating normally or when no power is present.

2

Power module LEDs.

- Steady green indicates that the module is connected to a power source.
- Steady red indicates that the module is not functioning within design specifications.
- Not illuminated indicates that there is no power to the module.

M 2002 Hardware testing from the CLI

You can use the Global **test hardware** command to test the hardware.

To use the **test hardware** command, you must establish an SSH connection and 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 output of the **test hardware** command is part of any generated error report.

The components are in the following categories.

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

Sample success statements are as follows.

- [success] Status of voltage reading 'Voltage PU +12' : ok.
- [success] CPUs OK
- [success] Memory all present
- [success] Statistics for interface 'eth10' show no errors
- [success] fan 1 operating within expected range

Sample 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)

Sample failure statements are as follows.

- [failure] Memory in error DIMM_H1, 0x001F
- [failure] fan 2 operating outside expected range (rpm too low)

M 2002 Running the diagnostic self-test

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

About this task

- Use the diagnostic self-test only when directed by IBM Support to help confirm a potential hardware problem with the appliance.
- The diagnostics user interface can differ depending on firmware release.

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 hear the fans start.
3. When you see DPOS boot - press <ESC> within 7 seconds for boot options, press ESC. You 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, 2020 - 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
- SKIP
- RUNNING
- SKIP
- n/a

M 2002 Status providers for sensors

The status providers for sensors can help diagnose hardware problems.

The appliance provides the following sensors status providers.

Fan speed sensors

Provides the measured speed in RPM for the fans in each fan module.

- In the GUI search field, enter `sensors` and click **Fan Sensors**.
- From the CLI, enter `show sensors-fans`.

Temperature sensors

Provides the measured temperature in degrees Celsius for internal components.

- In the GUI search field, enter `sensors` and click **Temperature Sensors**.
- From the CLI, enter `show sensors-temperature`.

Voltage sensors

Provides the measured voltage for internal components.

- In the GUI search field, enter `Sensors` and click **Voltage Sensors**.
- From the CLI, enter `show sensors-voltage`.

Current sensors

Provides the measured current for the internal components in milliamperes.

- In the GUI search field, enter `sensors` and click **Current Sensors**.
- From the CLI, enter `show sensors-current`.

RAID battery backup status

Monitors the power backup unit connected to the RAID controller.

- In the GUI search field, enter `RAID` and click **RAID Battery Backup Status**.
- From the CLI, enter `show raid-battery-module`.

Other sensors

Provides truth values for the status of intrusion switch and power supply modules, each hard disk in the array, and the battery.

- In the GUI search field, enter `sensors` and click **Other Sensors**.
- From the CLI, enter `show sensors-other`.

The values are as follows.

true

The indicator that the condition exists.

false

The indicator 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. M2002
 - Fault
 - Present

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.

About this task

To begin troubleshooting hardware problems, use this workflow to determine whether to contact IBM Support for assistance or to order a replacement part. The workflow guides you to the appropriate troubleshooting task.

Procedure

1. Did you receive a critical event through SNMP or 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.

Yes

Continue to step [“3” on page 151](#).

No

Continue to step [“2” on page 151](#).

2. Does the log file contain a critical message?

Yes

Continue to step [“3” on page 151](#).

No

Continue to step [“4” on page 151](#).

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.

- For a fan module, see [“Troubleshooting fan modules” on page 152](#).
- For a power supply module, see [“Troubleshooting power supply modules” on page 152](#)
- For a hard disk drive module, see [“Troubleshooting solid-state disk drive modules” on page 153](#).
- For a FRU, contact IBM Support.

No

Continue to step [“4” on page 151](#).

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

Yes

Continue with step [“5” on page 152](#).

No

See [“Troubleshooting the appliance” on page 153](#).

5. Are the LEDs lit for any modules?

Yes

- For a fan module, see [“Troubleshooting fan modules” on page 152.](#)
- For a power supply module, see [“Troubleshooting power supply modules” on page 152](#)
- For a hard disk drive module, see [“Troubleshooting solid-state disk drive modules” on page 153.](#)

No

See [“Troubleshooting the appliance” on page 153.](#)

M 2002 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, enter `show sensors -fans`.
- In the GUI search field, enter `sensors` and click **Fan Sensors**.
- 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.

If steady, the fan is operating at less than 1200 RPM or there is a fault.

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.

M 2002 Troubleshooting power supply modules

How to troubleshoot the power supply modules.

Procedure

1. View sensor status.

- From the CLI, run the `show sensors -other` command.
- In the GUI search field, enter `Sensors` and click **Other Sensors**.

2. Check whether the power supply model LED is red or flashing red or is not illuminated.

If red or flashing red, a problem could exist. If not illuminated, the module is not correctly connected to a power source.

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.

M 2002 Troubleshooting solid-state disk drive modules

How to troubleshoot the solid-state disk drive modules.

Procedure

1. View RAID status.

- From the GUI search field, enter RAID and click **RAID physical drive**.
- From the CLI, run the **show raid-physical-drive** command.

If the state shows Unconfigured Bad, the solid-state disk drive is damaged and must be replaced.

2. Contact IBM Support.

M 2002 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 diagnostic self-test to troubleshoot your appliance.

Related concepts

[“Hardware testing from the CLI” on page 148](#)

You can use the Global **test hardware** command to test the hardware.

Related tasks

[“Running the diagnostic self-test” on page 149](#)

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

M 2002 Removing or replacing the appliance or parts

Appliance parts can be removed or replaced under certain conditions.

The appliance includes customer replaceable unit (CRU) and field replaceable unit (FRU).

CRU

Replacement of a CRU is your responsibility.

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*.

M 2002 Installation 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.
- 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 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.

M 2002 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.

M 2002 Part listings

Access the parts list that is specific to your hardware.

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

M 2002 M2002 part lists

The 8441 appliance includes CRU parts and FRU parts.

- [“CRU parts” on page 154](#)
- [“FRU parts” on page 156](#)

CRU parts

Replacement of a CRU is your responsibility.

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

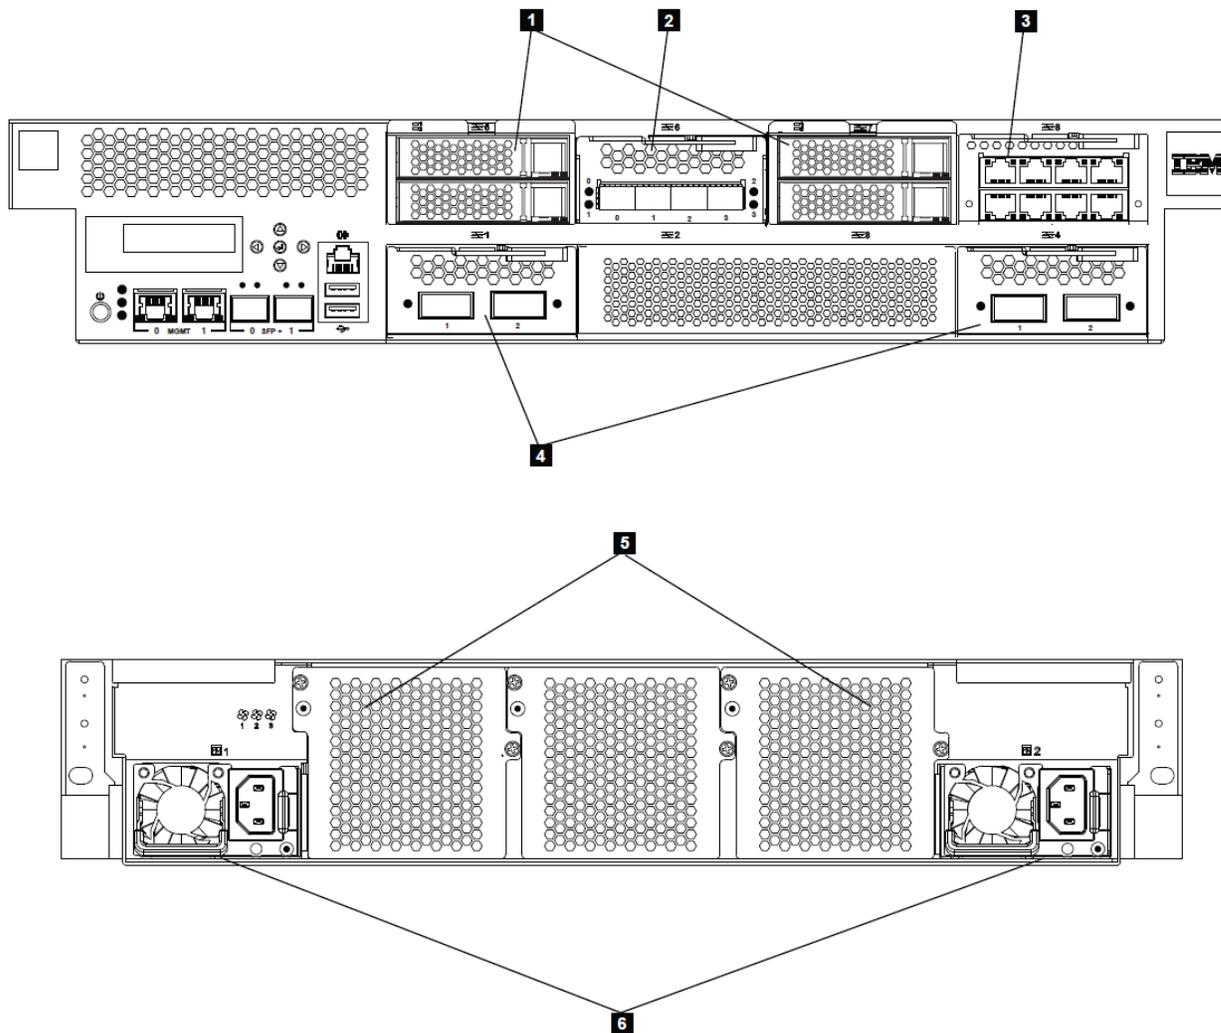


Figure 47. CRU part locations for the M2002 appliance

The following table lists the CRU parts.

Table 20. CRU part numbers for the M2002 appliance		
Index	Description	Part number
1	Solid-state disk drive module assembly (complete)	00VM462
2	10 Gb Ethernet module with 4 ports for SFP+ interface	01LK713
3	1 Gb Ethernet module with 8 ports for RJ45 interface	00VM052
4	40 Gb Ethernet with 2 ports for QSFP+ interface	01LK710
5	Fan module	01LK971
6	Power supply module	01LK969
The following CRU parts are not shown in the figure.		
	6 x SFP+ SR transceiver for 10 Gb ports	46N5368

Table 20. CRU part numbers for the M2002 appliance (continued)

Index	Description	Part number
	SFP+ 10GBE 5M cable	95Y1634
	4 x QSFP+ SR transceiver for 40 Gb ports	78P4902
	QSFP+ 40GBE 5M cable	00D5810
	DE-9 to RJ45 serial console cable	46N5656
	USB to RJ45 serial console cable	97Y0517
	2 x 1GBE Ethernet 10M cable	01AF038
	Rail kit to mount the appliance into the rack	60Y0328
	Appliance 8441-54X	01LK918

FRU parts

The following table lists the FRU parts. Replacement of a FRU must be performed by an IBM representative only.

Table 21. FRU part numbers for the M2002 appliance

Description	Part number
16 GB DDR4 DIMM	01LK977
32 GB M.2 SSD flash drive	01LK978
CMOS Button Cell battery	00RY543
RAID controller card	01LK973
RAID cache module, capacitor and cable	01LK974
SIM (Storage Interface Module)	01LK972

M 2002 Power cords

When you receive your appliance, the shipping carton contains power cords for connection to country-specific outlets and for rack-mounted appliances. Power cords are CRU parts.

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

Table 22. Power cord part numbers.

Country	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, 10A/250V 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

Table 22. Power cord part numbers. (continued)

Country	CRU part number	Description
Europe	39M5123	2.8m, 10A/250V, C13 to CEE 7/7
India	39M5226	2.8m, 10A/250V, C13 to IS 6538
Israel	39M5172	2.8m, 10A/250V, C13 to SI 32
Italy	39M5165	2.8m, 10A/240V, C13 to CEI 23-16
Japan	39M5186	2.8m, 12A/200V, C13 to JIS C-8303
	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
Switzerland	39M5158	2.8m, 10A/250V, C13 to SEV 1011-S24507
Taiwan	39M5247	2.8m, 10A/125V, C13 to CNS 10917-3
	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
	39M5095	2.8m, 10A/250V, C13 to NEMA 6-15P
Rack power cord (all countries)	39M5377	2.8m, 10A/125 - 250 VAC, IEC 320 C13 to IEC 320 C14

M 2002 Turning off the appliance

How to turn off power to the appliance.

Before you begin

Log in as an administrator or privileged user to save the running configuration to the persisted configuration.

- From the GUI, click **Save configuration**.
- From the CLI, use the **write memory** command.

Procedure

Gracefully shutdown of the appliance.

- From the GUI, perform the following procedure.
 - a. In the search field, enter **system**.
 - b. From the search results, click **System Control**.
 - c. Locate the **Shutdown** section.
 - d. From the **Mode** list, select **Power off System**.
 - e. In the **Delay** field, enter the amount of time to wait before shutdown process starts.
 - f. Click **Shutdown**.
- From the command line, run the **shutdown poweroff** command.
- From the appliance, press the power button.

What to do next

1. Verify that the power LED is not illuminated.
2. Unplug the power cords from both power supplies.

M 2002 Replacing a fan module

How to replace a failed fan module.

Before you begin

You must have part 01LK971 available.

About this task

The fan modules are hot-swap modules. When one or more fan modules are not working, replace the fan 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 any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)



DANGER: Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). (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. Remove the fan module.

The following figure shows the components that are mentioned.

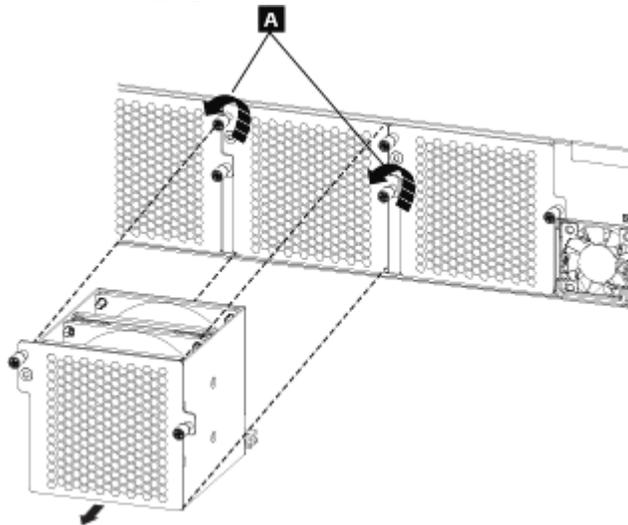


Figure 48. Removing a fan module

- a) Unscrew the two thumbscrews on the fan module until they twist without resistance **A**. 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. Set the faulty module aside.
 3. Unpack the replacement 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.

4. Carefully align the replacement module, and insert until the module face is flush with the rear panel.
5. Tighten the thumbscrews on the fan module.

What to do next

Confirm that the new module is working by verifying that the following are true.

- The fan module LED is not illuminated.
- The fault LED is not illuminated.

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

Related concepts

[“Fan modules” on page 133](#)

[“Returning an appliance or part” on page 172](#)

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

M 2002 Replacing a power supply module

How to replace a power supply module.

Before you begin

You must have part 01LK969 available.

About this task

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

- 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.



DANGER: Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)



DANGER: Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). (L002)

Procedure

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

The following figure shows the components that are mentioned.

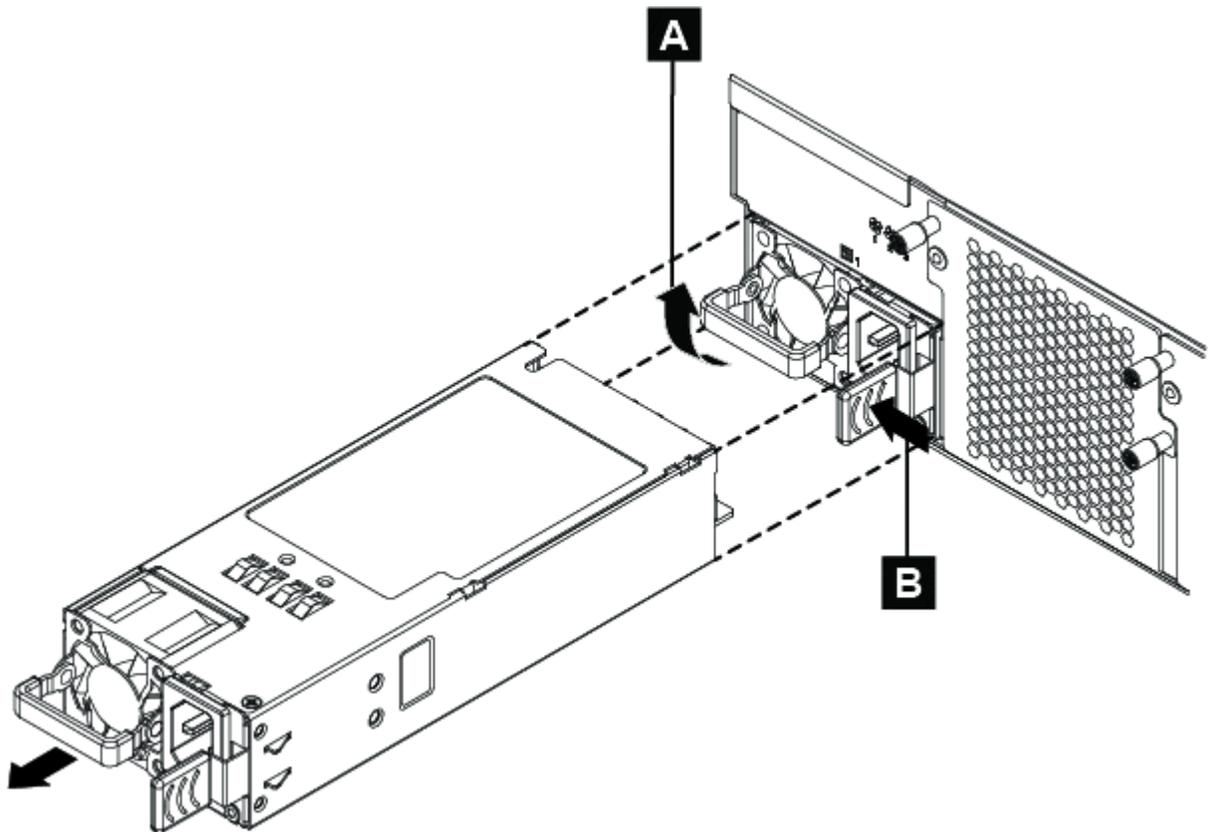


Figure 49. Removing a power supply module.

- a) Firmly grip the handle **A** of the failed module.
 - b) Push the release latch **B** toward the handle and hold in this position.
 - c) Pull the failed module from the appliance.
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.

Related concepts

[“Power supply modules” on page 134](#)

The appliance is powered by two redundant power supply modules.

[“Returning an appliance or part” on page 172](#)

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

M 2002 Replacing a solid-state disk drive module

How to replace the solid-state disk drive module.

Before you begin

You must have part 00VM462 available.

About this task

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

The modules are not hot-swap modules. Hot swapping modules can cause your system to crash, which might damage your appliance. You must turn off the appliance before you replace the module.



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:

- If IBM supplied power cords, 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.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening 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 the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
2. Attach all cables to the 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 turned on, complete a graceful shutdown by pressing the power button.
The power LED turns off indicates that the appliance is powered off.
2. Remove the module.
The following figure shows the numbered components that are mentioned.

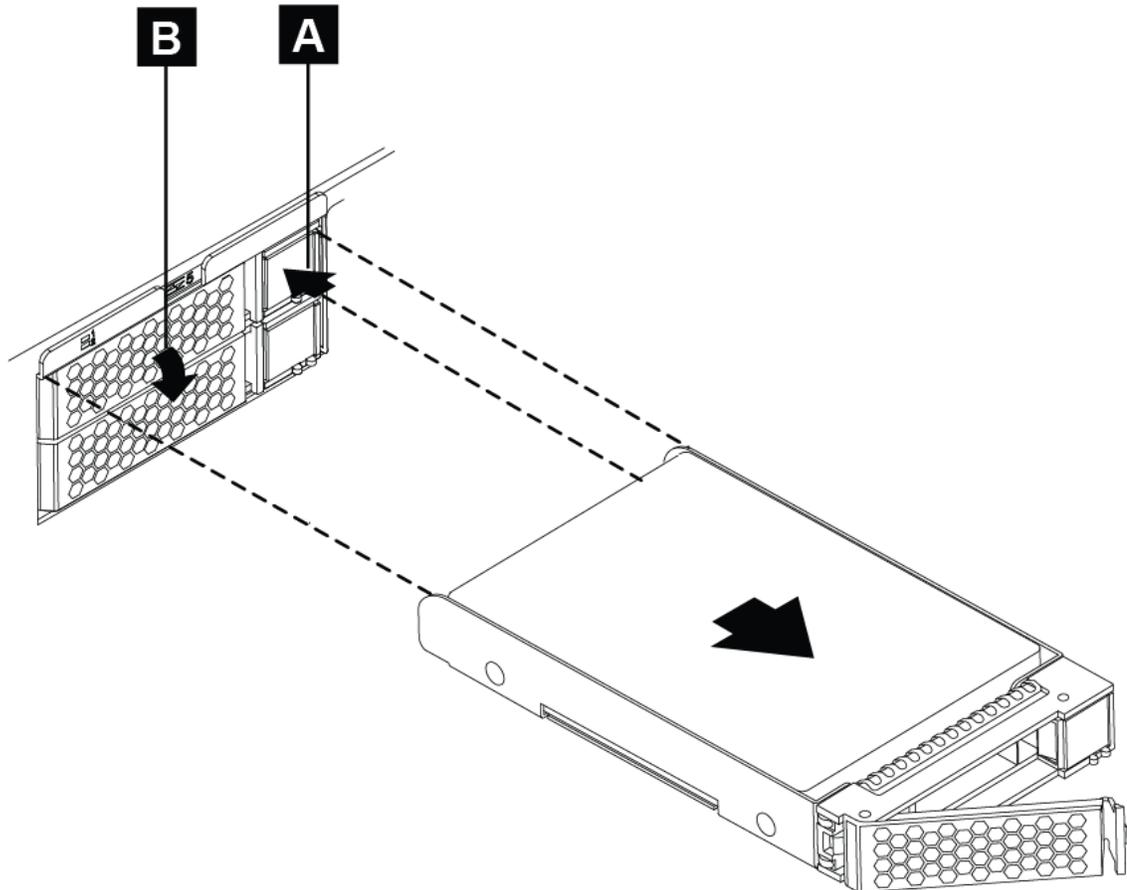


Figure 50. Removing a solid-state disk drive module.

- a) Press the locking arm release latch **A** and the locking arm is released.
 - b) To unlock the module, rotate the locking arm approximately 40 degrees by pulling out **B**.
 - c) Pull the module out of the appliance.
3. Set aside the failed module.
 4. Unpack the replacement 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.
 5. Insert the replacement module.
 - a) Carefully align the module, and insert into the opening until the module is seated.
 - b) Push the locking arm towards the appliance until the release latch clicks into place.
 6. Press the power button to turn on the appliance.

The power LED is illuminated.

What to do next

Verify that the new module is working.

- The solid-state disk drive activity LED illuminates.
- The solid-state disk state is not Unconfigured Bad.

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

Related concepts

[“Solid-state disk drive modules” on page 133](#)

The M2002 appliance has four solid-state disk drive modules.

[“Returning an appliance or part” on page 172](#)

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

Replacing an Ethernet module

How to replace an Ethernet module.

Before you begin

You must have the Ethernet module available.

- The part number of the 1 Gb Ethernet module is 00VM052.
- The part number of the 10 Gb Ethernet module is 01LK713.
- The part number of the 40 Gb Ethernet module is 01LK710.

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 all modules.

You can replace an Ethernet module if the following situation occurs or if directed by IBM Support.

- You cannot 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 GUI search field, enter ethernet and click **Ethernet Interfaces**.
 - 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:

- If IBM supplied power cords, 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.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening 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 the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
2. Attach all cables to the 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. When the power LED is no longer illuminated, the appliance is powered off.
2. Remove the Ethernet module.

The following figure shows the components that are mentioned.

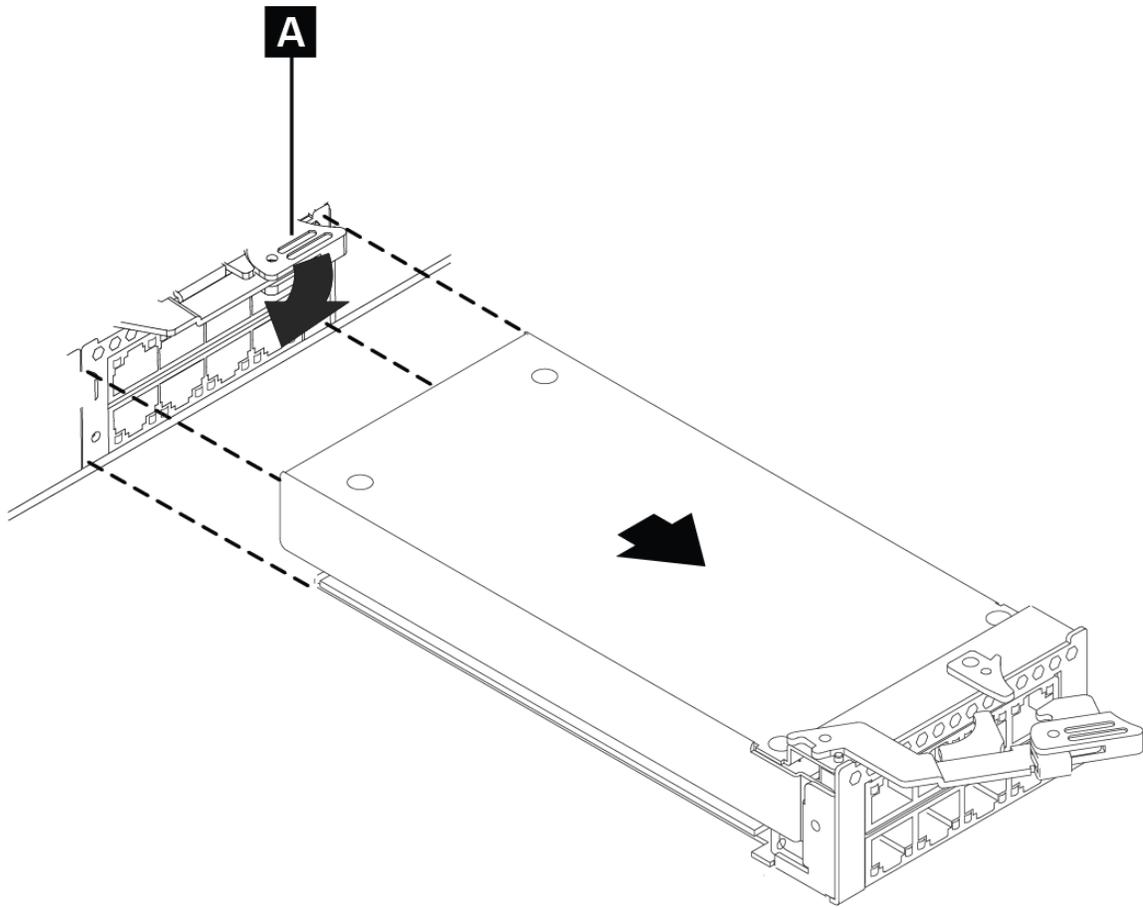


Figure 51. Removing a 1Gb Ethernet module.

- a) Grasp the blue latch **A**, rotate slightly, and pull outward.
- b) Pull the module completely out with care to support the module.
3. Set aside the Ethernet module.
4. Unpack the replacement 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. Carefully align the module, and insert into the appliance.
6. Push the Ethernet module forward until the module is securely in place.
7. Push the blue latch back in place to lock the module.
8. Turn on the appliance by pressing the power button, and verify that the power LED is steady green.

What to do next

Verify that the module is working, that you can connect to the network after you plug in the cable and the port activity LED is illuminated.

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

Related concepts

[“Network ports” on page 129](#)

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

[“Returning an appliance or part” on page 172](#)

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

Replacing an SFP+ or a QSFP+ transceiver

How to replace a 10 Gb SFP+ transceiver or a 40 Gb QSFP+ transceiver.

Before you begin

You must have part 46N5592 (10 Gb) or 78P4902 (40 Gb) available.

About this task

The appliance has six 10 Gb short reach transceivers, and four 40 Gb short reach transceivers. You remove a short reach transceiver to replace with a long reach transceiver or to use an SFP+ Copper Direct Attach twinaxial cable.



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:

- If IBM supplied power cords, 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.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening 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 the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
2. Attach all cables to the 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 power LED turns off indicates that the appliance is powered off.

2. Unplug the network cable from the transceiver.
3. Remove the SFP+ transceiver or the QSFP+ transceiver.

The following figure shows the one of the built-in 10 Gb SFP+ transceivers being removed.

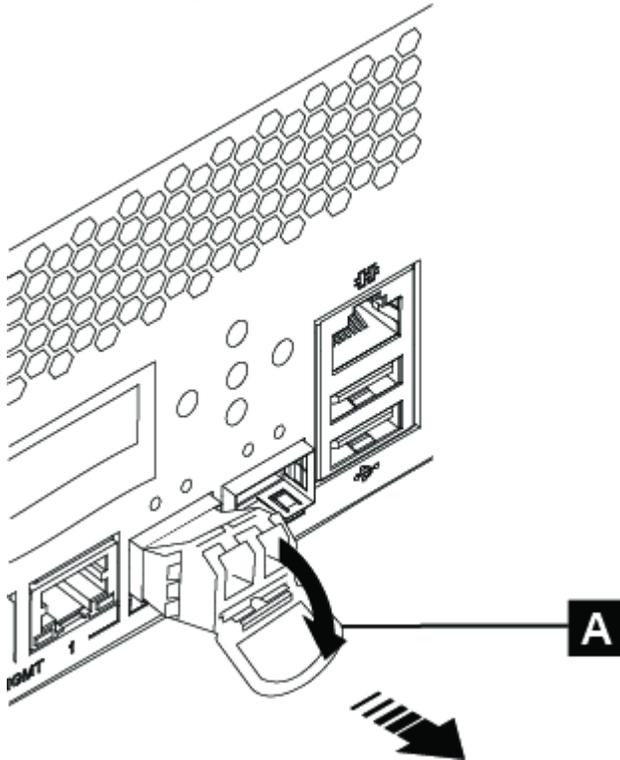
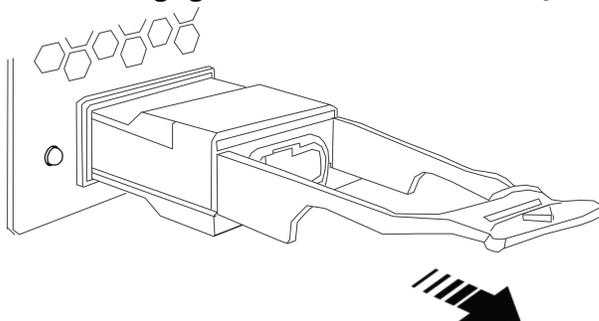


Figure 52. Removing the SFP+ transceiver

- a) Pull the release latch of the transceiver down **A**.
- b) Pull the transceiver out of the appliance by pulling on the release latch.

The following figure shows one of the 40 Gb QSFP+ transceivers being removed.



Pull the plastic tag to remove the transceiver

4. Set aside the transceiver.
5. Unpack the replacement unit.
6. Insert the replacement unit into the appliance until it is seated securely.
7. For a 10 Gb SFP+ transceiver, rotate the release latch upward into its locked position.

8. Reconnect any network cables.

M 2002 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. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). (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.
3. Separate the appliance from the rails.

The following figure shows the numbered components that are mentioned.

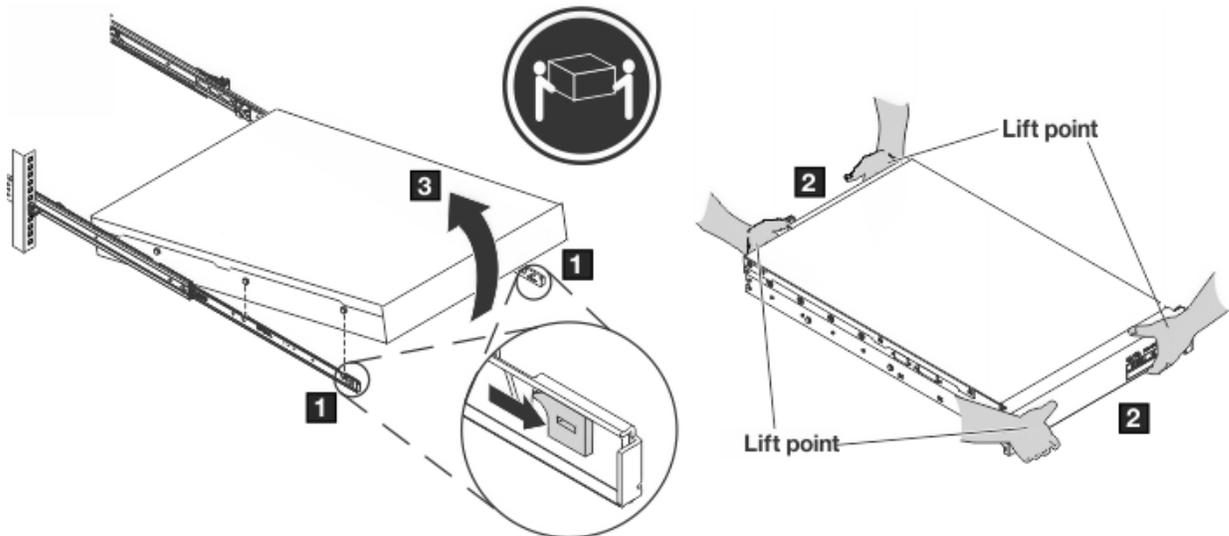


Figure 53. Unlatching and rotating the front of the appliance.

- a) Pull the locking levers **1** forward.
 - b) Make sure that two people support the front and the rear of the appliance at lift 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.
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 **2**.
5. Place the appliance on a sturdy, clean surface.
6. Slide the rails back in the rack.

Related concepts

“Returning an appliance or part” on page 172

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

M 2002 Replacing an appliance

How to replace the appliance.

Before you begin

You must have the following part available.

- The part number for an appliance (Model 8441-54X) is 01LK979.

You must back up your configuration.

About this task

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



DANGER: Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)



DANGER: Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). (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 turned on, complete a graceful shutdown by pressing the power button.

The power LED turns off indicates that the appliance is powered off.

2. Remove the appliance from the rack.

See [“Removing the appliance from the rack” on page 168](#).

3. Apply the repair identification (RID) tag.

When you receive a replacement appliance, the box contains an RID tag. The originally purchased appliance's serial number, also referred to as the "Entitlement Number" allows IBM to entitle the appliance for future maintenance or warranty service. Ensure that the Entitlement Number is recorded on the RID tag.

- a) Locate the defective appliance. Ensure that the serial number of the defective appliance matches the serial number that was reported to IBM. The serial number is on the label that is attached to the underside of the appliance and on a sticker at the left rail release latch.

- b) With a ballpoint pen, transcribe the machine type, model, and serial number of the defective appliance (the Entitlement Number) to the RID tag.
 - c) Secure the RID tag as close as possible, but not covering, the serial number of the replacement appliance. The serial number of the replacement appliance must be visible.
4. Install the new appliance on the slide rails and slide the appliance into the rack.
See [“Installing the appliance on the rails” on page 138](#).
 5. Reattach all network cables and power cords to their proper location.
 6. Reinitialize the appliance.

What to do next

After you verify that the new appliance is working, return the failed part to IBM.

If you have a backup and want to restore the appliance configuration, see the information in IBM Knowledge Center.

Related concepts

[“Returning an appliance or part” on page 172](#)

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

Related tasks

[“Setting up the initial firmware configuration” on page 142](#)

How to perform the base, initial firmware configuration.

Related information

IBM Knowledge Center: IBM DataPower Gateway (<http://www.ibm.com/support/knowledgecenter/SS9H2Y>)

Removing 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)

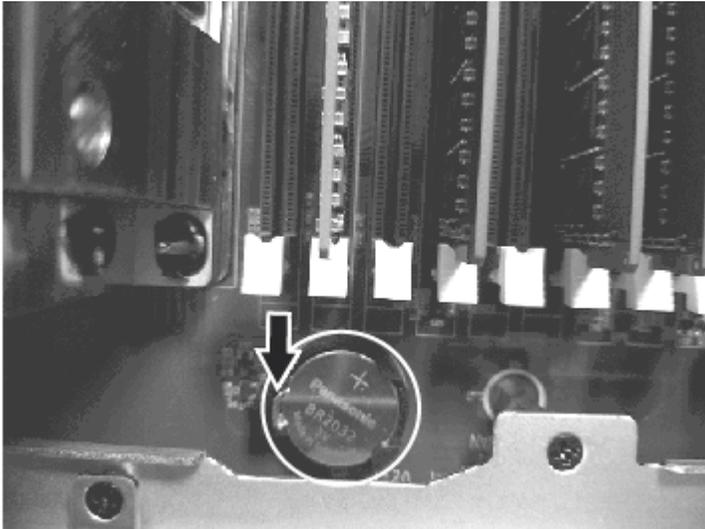


DANGER: Improper disposal or incineration of batteries or capacitors can cause life-threatening injury.

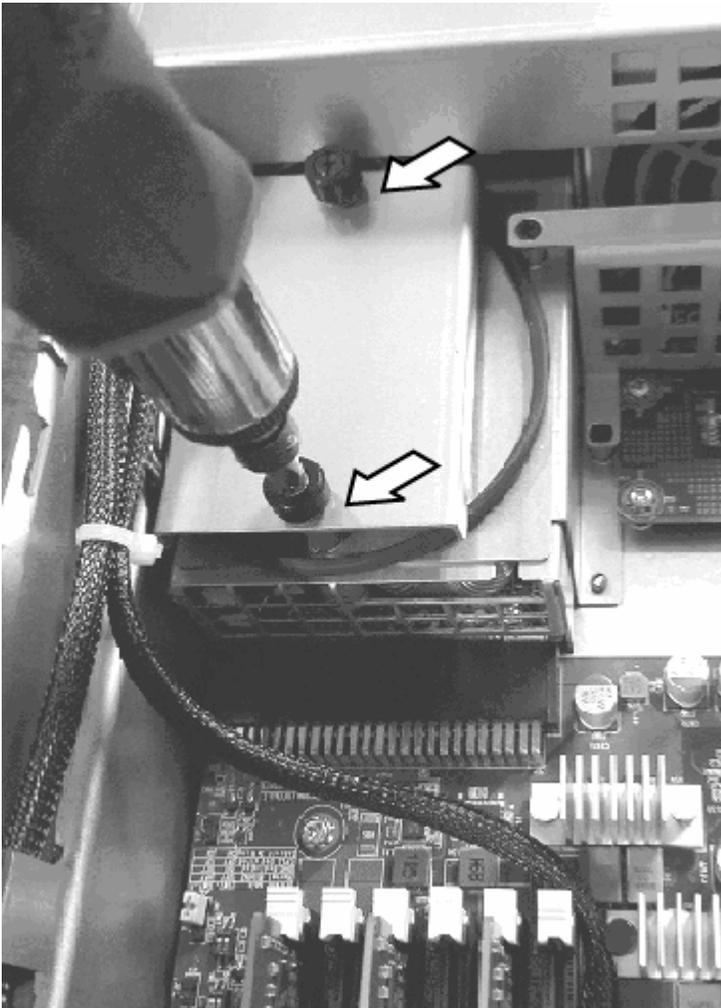
The appliance does not have any internal user serviceable parts. Any battery or capacitor is to be accessed and removed only by trained personnel. These instructions apply only to end-of-life recycling procedures.

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.
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 cache capacitor inside the chassis.



6. Loosen the indicated capacitor cover retention screws to remove the capacitor cover.
7. Disconnect the 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.

M 2002 **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.

Upgrading and downgrading firmware

To upgrade your IBM MQ Appliance, you 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 new firmware

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

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.

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 373 for guidance.

You can use the back up and restore capabilities to provide a recovery or rollback capability for your queue managers during the installation process. You can take an archive of your queue managers and data before you start your queue managers in the new environment. If you need to roll data back to the pre-upgrade state, the queue manager data can be restored from the archive file, and the appliance firmware rolled back to the last version. This use of archive files is useful for stand-alone (non-HA/DR) queue managers where you want to ensure that you have a backout strategy during migration. See [“Backing up a queue manager”](#) on page 365 for instructions on backing up and restoring.

You can install the new firmware either by using the IBM MQ Appliance web UI or by using the command line.

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.

Installing a new level of firmware by using the command line

After you have downloaded the new level of firmware, you can install it by using the command line.

Before you begin

If you are installing firmware on a high availability (HA) configuration, see the topic “Installing new firmware” on page 173 before you start installing. This topic gives guidance on upgrading one appliance at a time and so remaining operational.

Procedure

1. Back up your IBM MQ Appliance as described in [“Back up and restore” on page 358](#).
2. Restart the appliance as described in [“Restarting the appliance” on page 356](#).
3. Ensure that all the queue managers are stopped.
4. Copy the firmware image from the computer that you downloaded it to to the `image:` location on the appliance:
 - a) Connect to the command line of the appliance as described in [“Command line access” on page 207](#).
 - 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:
```

5. 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.

6. When the appliance restarts, verify that the firmware image is upgraded by entering the following commands:

```
show version
```

and, from the `mqcli` prompt:

```
dspmqver -v
```

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

[“boot image” on page 825](#)

This command designates the startup firmware image and restarts the appliance with this image.

Installing a new level of firmware by using the IBM MQ Appliance web UI

After you have downloaded the new level of firmware, you can install it by using the IBM MQ Appliance web UI.

Before you begin

If you are installing firmware on a high availability (HA) configuration, see the topic “Installing new firmware” on page 173 before you start installing. This topic gives guidance on upgrading one appliance at a time and so remaining operational.

Procedure

1. Back up your IBM MQ Appliance as described in [“Back up and restore”](#) on page 358.
2. Start the web UI as described in [“Configuring the IBM MQ Appliance web UI”](#) on page 209.
3. Restart the appliance:
 - a) Click the administration icon  and select **Main > System Control**.
 - b) Set the Shutdown **Mode** to **Reboot system**.
 - c) Click **Shutdown**.
4. Click the console icon  to switch to the IBM MQ Console and ensure that all the queue managers are stopped.
5. Copy the new firmware image to the appliance and restart the system:
 - a) Click the administration icon  and select **Main > System Control**.
 - b) In the **Boot Image** section, click **Upload**, and browse your local system for the new firmware image in the upload file window. Click **Upload** to copy the file to the appliance.
 - c) Select **I accept the terms of the license agreement**.
 - d) Click **Boot Image**.
6. When the appliance restarts, verify that the firmware image is upgraded:
 - a) Click the status icon  and select **System > Version Information**.
 - b) Check that the version information is as you expect.

Reverting firmware

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

When you upgrade the IBM MQ Appliance firmware, 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 perform the **boot switch** operation, all system configuration reverts to its state at initial upgrade. 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. If you delete a high availability group and then later perform a **boot switch** operation, the deleted group will reappear after the reversion (if it existed at the earlier version). You must delete the reappeared group before proceeding.

Reverting to the previous level of firmware by using the command line

You can revert to a previous level of firmware by using the command line.

Procedure

1. Connect to the command line of the appliance as described in [“Command line access” on page 207](#).
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
```

Related reference

[“boot switch” on page 826](#)

This command switches between primary and secondary installations.

Reverting to the previous level of firmware by using the IBM MQ Appliance web UI

You can revert to a previous level of firmware by using the IBM MQ Appliance web UI.

Procedure

1. Start the web UI as described in [“Configuring the IBM MQ Appliance web UI” on page 209](#).
2. Click the administration icon  and select **Main > System Control**.
3. Click **Switch Installation Image**.
4. When prompted, confirm the switch boot image operation.

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 firmware” on page 175](#).

To install a specific firmware version, follow the procedure described in [“Installing new firmware” on page 173](#).

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 358](#).
2. Restart the appliance by following the procedure described in [“Restarting the appliance” on page 356](#).

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.

V 9.1.0 Upgrading a version 8.0 IBM MQ Appliance to version 9.1

You can upgrade a version 8.0 IBM MQ Appliance (running version 8.0 queue managers) to version 9.1 (running version 9.1 queue managers).

There are a number of options in planning your upgrade, and the best approach depends on your current environment, disaster recovery strategy, and planned continuity of service during the upgrade window.

If you have a high availability (HA) configuration (with or without disaster recovery), you must follow the procedure to update your HA queue managers, see [“Upgrading a Version 8.0 HA configuration to version 9.1” on page 180](#) or [“Upgrading a Version 8.0 HA configuration with DR to version 9.1” on page 181](#).

If you currently have a disaster recovery (DR) configuration, you can keep queue managers active through most of the upgrade process. You do this by running a queue manager on one appliance in the DR pair until the upgrade is complete, then swapping appliances and repeating the upgrade. Note that, once a queue manager has been started on an upgraded appliance, this queue manager will not be able to fail back to an appliance running the earlier (Version 8.0) firmware. For this reason, you should upgrade DR recovery appliances before live appliances.

The Version 9.1 firmware includes new capabilities for simple backup and restore of queue manager data, see [“Backing up a queue manager” on page 365](#). These capabilities can be used to provide a recovery or rollback capability for your Version 8 queue managers during the upgrade process. You can take an archive of your queue managers and data before you start your queue managers in the new environment. If you need to roll data back to the pre-upgrade state, the queue manager data can be restored from the archive file, and the appliance firmware rolled back to the last version. This use of archive files is useful for stand-alone (non-HA/DR) queue managers where you want to ensure that you have a backout strategy during migration.

The procedure to upgrade a stand-alone appliance using the described strategy contains several major steps:

1. Stop all your queue managers.
2. Upgrade your appliance to the new firmware level.
3. Create the location and allocate storage for your backup archive files.
4. Back up all of your queue managers to archive files.
5. Restart your queue managers.

You can complete these steps by using the appliance command line interface. You can also complete most of the steps by using the IBM MQ Console.

The IBM MQ Console has completely changed for version 9.0. You cannot migrate your existing console layout from your version 8.0 appliance, but you should take a backup of your console layout before you upgrade in case you want to subsequently revert to version 8.0. See [Backing up IBM MQ Appliance web UI configuration data](#) in the IBM MQ Appliance version 8.0 documentation for instructions.

User authentication and authorization has completely changed for version 9.0 and is now achieved through role based management (RBM). If you have controlled user authorization by setting up user groups with limited access to certain appliance command groups in version 8.0, then you will need to manually reconstruct these authorizations by using RBM (see [“Credential mapping with local user groups”](#) on page 464). Be sure to make a note of current user group command groups before you upgrade because these groups are not visible after you upgrade.

V 9.1.0 Upgrading to version 9.1 by using the command line

You can use the command line interface to upgrade your version 8.0 IBM MQ Appliance to version 9.1.

Before you begin

If you are upgrading a disaster recovery (DR) configuration, see the topic [“Upgrading a version 8.0 IBM MQ Appliance to version 9.1”](#) on page 177 before you start upgrading. This topic gives guidance on upgrading one appliance at a time and so remaining operational.

If you are upgrading a high availability (HA) configuration, see the topics [“Upgrading a Version 8.0 HA configuration to version 9.1”](#) on page 180 and [“Upgrading a Version 8.0 HA configuration with DR to version 9.1”](#) on page 181.

About this task

Follow this procedure to take back ups of your queue managers, and associated logs and data, and then upgrade your appliance to IBM MQ version 9.1.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.
2. Log in as a user in the administrators group.
3. Type **mqc1i** to enter IBM MQ configuration mode.
4. Stop each of your queue managers:

```
endmqm queue_manager_name
```

5. Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file). The appliance restarts.

6. Run a command to create the directory and allocate storage for your backup archive files:

```
createbackupfs -s size
```

Where *size* specifies the space that is allocated on the appliance RAID volume for back ups in GB. A directory is created at the location mqbackup:///QMgrs.

7. Back up each of your queue managers:

```
mqbackup -m queue_manager_name
```

Where *queue_manager_name* specifies the queue manager to be backed up. The archive file that is created is named *queue_manager_name.bak* by default.

8. Restart your queue managers:

```
stirmqm queue_manager_name
```

Where *queue_manager_name* is the queue manager to restart.

V 9.1.0 Upgrading to version 9.1 by using the IBM MQ Console

You can use the IBM MQ Console to upgrade your version 8.0 IBM MQ Appliance to version 9.1.

Before you begin

If you are upgrading a disaster recovery (DR) configuration, see the topic [“Upgrading a version 8.0 IBM MQ Appliance to version 9.1” on page 177](#) before you start upgrading. This topic gives guidance on upgrading one appliance at a time and so remaining operational.

If you are upgrading a high availability (HA) configuration, see the topics [“Upgrading a Version 8.0 HA configuration to version 9.1” on page 180](#) and [“Upgrading a Version 8.0 HA configuration with DR to version 9.1” on page 181](#).

About this task

Follow this procedure to take back ups of your queue managers, and associated logs and data, and then your appliance to IBM MQ version 9.1.

Procedure

1. Start the IBM MQ Console as described in [“Using the IBM MQ Console” on page 311](#).
2. Stop each of your queue managers:
 - a) Select the queue manager that you want to stop from the list in the local queue manager widget.



- b) Click the stop icon in the local queue manager widget toolbar.
 - c) Confirm that you want to stop the queue manager by clicking **Stop**.
3. Upgrade your appliance to the new firmware level. For this operation you use the IBM MQ Appliance web UI:
 - a) Start the IBM MQ Appliance web UI, and click the administration icon .
 - b) Select **Main > System Control**
 - c) In the **Boot Image** section, select the license accept option. Select the image that you want to use from the **Firmware File** list. This list contains all the files in the Image directory on the appliance. (You can use the buttons in this window to upload files to the appliance or move them to the Image directory as required.)
 - d) Click **Boot Image** to restart the appliance with the new firmware installed.
 4. When the appliance has restarted, create the location and allocate storage for your backup archive files:

```
createbackupfs -s size
```

Where *size* specifies the space that is allocated on the appliance RAID volume for back ups in GB. A directory is created at the location `mqbackup:///QMgrs`.

5. Open the IBM MQ Console again. Back up each of your queue managers:
 - a) Select the queue manager that you want to back up from the list in the local queue manager widget.
 - b) Select **Back Up** from the widget menu.
6. Restart your queue managers:
 - a) Select the queue manager that you want to restart from the list in the local queue manager widget.



- b) Click the start icon in the local queue manager widget toolbar.

V 9.1.0 Upgrading a Version 8.0 HA configuration to version 9.1

You can upgrade your high availability (HA) configuration while maintaining the availability of your HA queue managers.

Before you begin

Before upgrading, you should back up your queue managers as described in “[Backing up a queue manager](#)” on page 365. You should also ensure that none of your queue managers are in the partitioned state, resolving the state before upgrade if necessary.

About this task

You must use the command line interface to upgrade your HA queue managers from version 8.0 to version 9.1.

You suspend one appliance so that any queue managers running there fail over to the other appliance in your HA pair. You then upgrade the firmware on the first appliance and restart it. After restart you run a command to upgrade the queue managers. Finally you upgrade the second HA appliance.

The additional manual step of upgrading the queue managers is required to configure the HA group in preparation for the change to Version 9.1. When an appliance upgrades to Version 9.1, you see a new status of `migration pending` for your HA queue managers. You must then run a `migrate` command to move the queue managers to the upgraded appliance and to the new version (note that you cannot use the `sethapreferred` command to move the queue manager to the upgraded appliance at this point). After migration, these queue managers cannot fail back to the other appliance until you upgrade that appliance to Version 9.1 as well. For this reason, you might want to start by migrating a test queue manager. If you then decide, for some reason, to move back to the previous firmware installation, you can do so and resume all the other, non-migrated queue managers.

Procedure

- To upgrade the first appliance in your HA group:
 - a) Suspend the appliance from the HA group:

```
sethagrp -s
```

Queue managers fail over to the other appliance.

- b) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see “[Installing a new level of firmware by using the command line](#)” on page 174 for details on obtaining the file).

The appliance restarts. The queue managers in the HA group show the status `Migration Pending`.

Note: At this point, multiple messages are logged in `mqtrace:messages` on both appliances indicating that Pacemaker cannot communicate. You should upgrade the second appliance as soon as possible and remember to perform housekeeping on `mqtrace:messages` afterwards.

- c) Optionally, back up your queue managers as described in “[Backing up a queue manager](#)” on page 365. You can then, if required, back out of your upgrade by reverting to the previous firmware version and restoring your queue managers.
- d) Upgrade each queue manager by entering the following command:

```
sethagrp -u QMName
```

Each queue manager now switches over to this appliance. The other appliance in the HA group must be available, or the command fails. After the queue managers have switched over, they have

the ended status on the second appliance and cannot fail back to the second appliance until that too is upgraded.

- To upgrade the second appliance:
 - a) Upgrade the appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

The appliance restarts.

Related reference

[“sethagr \(set HA group\)”](#) on page 675

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“boot image”](#) on page 825

This command designates the startup firmware image and restarts the appliance with this image.

V 9.1.0 Upgrading a Version 8.0 HA configuration with DR to version 9.1

You can upgrade your high availability (HA) configuration while maintaining the availability of your HA queue managers. You upgrade the disaster recovery (DR) node first.

Before you begin

Before upgrading, you should back up your queue managers as described in [“Backing up a queue manager”](#) on page 365. You should also ensure that none of your queue managers are in the partitioned state, resolving the state before upgrade if necessary.

About this task

You must use the command line interface to upgrade your DR node and upgrade your HA queue managers from version 8.0 to version 9.1.

First you upgrade your DR node. This is to ensure that queue manager data does not run on an upgraded appliance, unable to fail over in the event of a disaster to an appliance that has not yet been upgraded. Then you upgrade the HA pair. You suspend one appliance so that any queue managers running there fail over to the other appliance in your HA pair. You then upgrade the firmware on the first appliance and restart it. After restart you run a command to upgrade the queue managers. Finally you upgrade the second HA appliance.

The additional manual step of upgrading the queue managers is required to configure the HA group in preparation for the change to Version 9.1. When an appliance upgrades to Version 9.1, you see a new status of `migration pending` for your HA queue managers. You must then run a `migrate` command to move the queue managers to the upgraded appliance and to the new version (note that you cannot use the **sethappreferred** command to move the queue manager to the upgraded appliance at this point). After migration, these queue managers cannot fail back to the other appliance until you upgrade that appliance to Version 9.1 as well. For this reason, you might want to start by migrating a test queue manager. If you then decide, for some reason, to move back to the previous firmware installation, you can do so and resume all the other, non-migrated queue managers.

Procedure

- To upgrade the DR node:
 - a) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

- b) Restart the appliance. Do not start any of the queue managers on the DR node.
- To upgrade the first appliance in your HA group:
 - a) Suspend the appliance from the HA group:

```
sethagrp -s
```

Queue managers fail over to the other appliance.

- b) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

- c) Restart the appliance. The queue managers in the HA group show the status as Migration Pending.

Note: At this point, multiple messages are logged in `mqtrace:messages` on both appliances indicating that Pacemaker cannot communicate. You should upgrade the second appliance as soon as possible and remember to perform housekeeping on `mqtrace:messages` afterwards.

- d) Optionally, back up your queue managers as described in [“Backing up a queue manager”](#) on page 365. You can then, if required, back out of your upgrade by reverting to the previous firmware version and restoring your queue managers.
- e) Upgrade each queue manager by entering the following command:

```
sethagrp -u QMName
```

Each queue manager now switches over to this appliance. The other appliance in the HA group must be available, or the command fails. After the queue managers have switched over, they have the ended status on the second appliance and cannot fail back to the second appliance until that too is upgraded.

- To upgrade the second appliance:
 - a) Upgrade the appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

- b) Restart the appliance.

Related reference

[“sethagrp \(set HA group\)”](#) on page 675

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“boot image”](#) on page 825

This command designates the startup firmware image and restarts the appliance with this image.

V 9.1.0 Upgrading a Version 8.0 DR configuration to version 9.1

You can upgrade your disaster recovery (DR) configuration, you upgrade the recovery node first.

Before you begin

Before upgrading, you should back up your queue managers as described in [“Backing up a queue manager”](#) on page 365. You should also ensure that none of your queue managers are in the partitioned state, resolving the state before upgrade if necessary.

About this task

You must use the command line interface to upgrade your DR configuration from version 8.0 to version 9.1.

You upgrade your recovery node first. This is to ensure that queue manager data does not run on an upgraded appliance, unable to fail over to an appliance that has not yet been upgraded.

Procedure

- To upgrade the recovery node:
 - a) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

- b) The appliance restarts. Do not start any of the queue managers on the recovery node.
- To upgrade the main appliance in your DR configuration:
 - a) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

The appliance restarts.

Related reference

[“boot image”](#) on page 825

This command designates the startup firmware image and restarts the appliance with this image.

V 9.1.0 Upgrading a version 9.0 IBM MQ Appliance to version 9.1

You can upgrade a version 9.0 IBM MQ Appliance (running version 9.0 queue managers) to version 9.1 (running version 9.1 queue managers).

There are a number of options in planning your upgrade, and the best approach depends on your current environment, disaster recovery strategy, and planned continuity of service during the upgrade window.

If you have a high availability (HA) configuration (with or without disaster recovery), you must follow the procedure to update your HA queue managers, see [“Upgrading a Version 9.0 HA configuration to version 9.1”](#) on page 186 or [“Upgrading a Version 9.0 HA configuration with DR to version 9.1”](#) on page 187.

If you currently have a disaster recovery (DR) configuration, you can keep queue managers active throughout the upgrade process. You do this by running a queue manager on one appliance in the DR pair until the upgrade is complete, then swapping appliances and repeating the upgrade. Note that, once a queue manager has been started on an upgraded appliance, this queue manager will not be able to fail back to an appliance running the earlier (Version 9.0) firmware. For this reason, you should upgrade DR recovery appliances before live appliances.

You can take an archive of your queue managers and data before you start your queue managers in the new environment. If you need to roll data back to the pre-upgrade state, the queue manager data can be restored from the archive file, and the appliance firmware rolled back to the last version. This use of archive files is useful for stand-alone (non-HA/DR) queue managers where you want to ensure that you have a backout strategy during migration.

The procedure to upgrade a stand-alone appliance using the described strategy contains several major steps:

1. Stop all your queue managers.
2. Create the location and allocate storage for your backup archive files.

3. Back up all of your queue managers to archive files.
4. Upgrade your appliance to the new firmware level.
5. Restart your queue managers.

You can complete these steps by using the appliance command line interface. You can also complete most of the steps by using the IBM MQ Console.

▶ V 9.1.0 Upgrading to version 9.1 by using the command line

You can use the command line interface to upgrade your version 9.0 IBM MQ Appliance to version 9.1.

Before you begin

If you are upgrading a high availability (HA) configuration, see the instructions in [“Upgrading a Version 9.0 HA configuration to version 9.1” on page 186](#). If you are upgrading and HA configuration that also has disaster recovery, see [“Upgrading a Version 9.0 HA configuration with DR to version 9.1” on page 187](#).

About this task

Follow this procedure to take back ups of your queue managers, and associated logs and data, and then upgrade your appliance to IBM MQ version 9.1.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#).
2. Log in as a user in the administrators group.
3. Type **mqcli** to enter IBM MQ configuration mode.
4. Stop each of your queue managers:

```
endmqm queue_manager_name
```

5. Run a command to create the directory and allocate storage for your backup archive files:

```
createbackupfs -s size
```

Where *size* specifies the space that is allocated on the appliance RAID volume for back ups in GB. A directory is created at the location `mqbackup:///QMGRS`.

6. Back up each of your queue managers:

```
mqbackup -m queue_manager_name
```

Where *queue_manager_name* specifies the queue manager to be backed up. The archive file that is created is named *queue_manager_name.bak* by default.

7. Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line” on page 174](#) for details on obtaining the file). The appliance restarts.

8. Restart your queue managers:

```
strmqm queue_manager_name
```

Where *queue_manager_name* is the queue manager to restart.

V 9.1.0 Upgrading to version 9.1 by using the IBM MQ Console

You can use the IBM MQ Console to upgrade your version 8.0 IBM MQ Appliance to version 9.0.

Before you begin

If you are upgrading a high availability (HA) configuration, see the instructions in “[Upgrading a Version 9.0 HA configuration to version 9.1](#)” on page 186. If you are upgrading and HA configuration that also has disaster recovery, see “[Upgrading a Version 9.0 HA configuration with DR to version 9.1](#)” on page 187.

About this task

Follow this procedure to take back ups of your queue managers, and associated logs and data, and then your appliance to IBM MQ version 9.1.

Procedure

1. Start the IBM MQ Console as described in “[Using the IBM MQ Console](#)” on page 311.
2. Stop each of your queue managers:
 - a) Select the queue manager that you want to stop from the list in the local queue manager widget.

A blue rectangular button with the word "Stop" in white text and a white square icon with a circle inside.

- b) Click the stop icon in the local queue manager widget toolbar.
 - c) Confirm that you want to stop the queue manager by clicking **Stop**.
3. Create the location and allocate storage for your backup archive files:

```
createbackupfs -s size
```

Where *size* specifies the space that is allocated on the appliance RAID volume for back ups in GB. A directory is created at the location `mqbackup:///QMGRS`.

4. Open the IBM MQ Console again. Back up each of your queue managers:
 - a) Select the queue manager that you want to back up from the list in the local queue manager widget.
 - b) Select **Back Up** from the widget menu.
5. Upgrade your appliance to the new firmware level. For this operation you use the IBM MQ Appliance web UI:



- a) Start the IBM MQ Appliance web UI, and click the administration icon.
 - b) Select **Main > System Control**
 - c) In the **Boot Image** section, select the license accept option. Select the image that you want to use from the **Firmware File** list. This list contains all the files in the Image directory on the appliance. (You can use the buttons in this window to upload files to the appliance or move them to the Image directory as required.)
 - d) Click **Boot Image** to restart the appliance with the new firmware installed.
6. Restart your queue managers:
 - a) Select the queue manager that you want to restart from the list in the local queue manager widget.

A blue rectangular button with the word "Start" in white text and a white play button icon.

- b) Click the start icon in the local queue manager widget toolbar.

V 9.1.0 Upgrading a Version 9.0 HA configuration to version 9.1

You can upgrade your high availability (HA) configuration while maintaining the availability of your HA queue managers.

Before you begin

Before upgrading, you should back up your queue managers as described in “[Backing up a queue manager](#)” on page 365. You should also ensure that none of your queue managers are in the partitioned state, resolving the state before upgrade if necessary.

About this task

You must use the command line interface to upgrade your HA queue managers from version 9.0 to version 9.1.

You suspend one appliance so that any queue managers running there fail over to the other appliance in your HA pair. You then upgrade the firmware on the first appliance and restart it. After restart you run a command to upgrade the queue managers. Finally you upgrade the second HA appliance.

The additional manual step of upgrading the queue managers is required to configure the HA group in preparation for the change to Version 9.1. When an appliance upgrades to Version 9.1, you see a new status of `migration pending` for your HA queue managers. You must then run a `migrate` command to move the queue managers to the upgraded appliance and to the new version (note that you cannot use the **`sethappreferred`** command to move the queue manager to the upgraded appliance at this point). After migration, these queue managers cannot fail back to the other appliance until you upgrade that appliance to Version 9.1 as well. For this reason, you might want to start by migrating a test queue manager. If you then decide, for some reason, to move back to the previous firmware installation, you can do so and resume all the other, non-migrated queue managers.

Procedure

- To upgrade the first appliance in your HA group:
 - a) Suspend the appliance from the HA group:

```
sethagrp -s
```

Queue managers fail over to the other appliance.

- b) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see “[Installing a new level of firmware by using the command line](#)” on page 174 for details on obtaining the file).

The appliance restarts. The queue managers in the HA group show the status `Migration Pending`.

Note: At this point, multiple messages are logged in `mqtrace:messages` on both appliances indicating that Pacemaker cannot communicate. You should upgrade the second appliance as soon as possible and remember to perform housekeeping on `mqtrace:messages` afterwards.

- c) Upgrade each queue manager by entering the following command:

```
sethagrp -u QMName
```

Each queue manager now switches over to this appliance. The other appliance in the HA group must be available, or the command fails. After the queue managers have switched over, they have the `ended` status on the second appliance and cannot fail back to the second appliance until that too is upgraded.

- To upgrade the second appliance:
 - a) Upgrade the appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

The appliance restarts.

Related reference

[“sethagr \(set HA group\)”](#) on page 675

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“boot image”](#) on page 825

This command designates the startup firmware image and restarts the appliance with this image.

V 9.1.0 Upgrading a Version 9.0 HA configuration with DR to version 9.1

You can upgrade your high availability (HA) configuration while maintaining the availability of your HA queue managers. You upgrade the disaster recovery (DR) node first.

Before you begin

Before upgrading, you should back up your queue managers as described in [“Backing up a queue manager”](#) on page 365. You should also ensure that none of your queue managers are in the partitioned state, resolving the state before upgrade if necessary.

About this task

You must use the command line interface to upgrade your DR node and upgrade your HA queue managers from version 9.0 to version 9.1.

First you upgrade your DR node. This is to ensure that queue manager data does not run on an upgraded appliance, unable to fail over in the event of a disaster to an appliance that has not yet been upgraded. Then you upgrade the HA pair. You suspend one appliance so that any queue managers running there fail over to the other appliance in your HA pair. You then upgrade the firmware on the first appliance and restart it. After restart you run a command to upgrade the queue managers. Finally you upgrade the second HA appliance.

The additional manual step of upgrading the queue managers is required to configure the HA group in preparation for the change to Version 9.1. When an appliance upgrades to Version 9.1, you see a new status of `migration pending` for your HA queue managers. You must then run a `migrate` command to move the queue managers to the upgraded appliance and to the new version (note that you cannot use the `sethappreferred` command to move the queue manager to the upgraded appliance at this point). After migration, these queue managers cannot fail back to the other appliance until you upgrade that appliance to Version 9.1 as well. For this reason, you might want to start by migrating a test queue manager. If you then decide, for some reason, to move back to the previous firmware installation, you can do so and resume all the other, non-migrated queue managers.

Procedure

- To upgrade the DR node:
 - a) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

- b) The appliance restarts. Do not start any of the queue managers on the DR node.
- To upgrade the first appliance in your HA group:
 - a) Suspend the appliance from the HA group:

```
sethagrp -s
```

Queue managers fail over to the other appliance.

- b) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

The appliance restarts. The queue managers in the HA group show the status as Migration Pending.

Note: At this point, multiple messages are logged in `mqtrace:messages` on both appliances indicating that Pacemaker cannot communicate. You should upgrade the second appliance as soon as possible and remember to perform housekeeping on `mqtrace:messages` afterwards.

- c) Upgrade each queue manager by entering the following command:

```
sethagrp -u -m QMName
```

Each queue manager now switches over to this appliance. The other appliance in the HA group must be available, or the command fails. After the queue managers have switched over, they have the ended status on the second appliance and cannot fail back to the second appliance until that too is upgraded.

- To upgrade the second appliance:
 - a) Upgrade the appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

- b) Restart the appliance.

Related reference

[“sethagrp \(set HA group\)”](#) on page 675

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“boot image”](#) on page 825

This command designates the startup firmware image and restarts the appliance with this image.

V 9.1.0 Upgrading a Version 9.0 DR configuration to version 9.1

You can upgrade your disaster recovery (DR) configuration, you upgrade the recovery node first.

Before you begin

Before upgrading, you should back up your queue managers as described in [“Backing up a queue manager”](#) on page 365. You should also ensure that none of your queue managers are in the partitioned state, resolving the state before upgrade if necessary.

About this task

You must use the command line interface to upgrade your DR configuration from version 9.0 to version 9.1.

You upgrade your recovery node first. This is to ensure that queue manager data does not run on an upgraded appliance, unable to fail over to an appliance that has not yet been upgraded.

Procedure

- To upgrade the recovery node:
 - a) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

- b) The appliance restarts. Do not start any of the queue managers on the recovery node.
- To upgrade the main appliance in your DR configuration:
 - a) Upgrade your appliance to the new firmware level:

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

Where *firmware_file* is the file that contains the appliance firmware upgrade (see [“Installing a new level of firmware by using the command line”](#) on page 174 for details on obtaining the file).

The appliance restarts.

Related reference

[“boot image”](#) on page 825

This command designates the startup firmware image and restarts the appliance with this image.

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.

You can transfer queue managers from older appliances to an M2002 appliance. You can upgrade the queue managers to a later release as you do this transfer.

The M2002 appliance runs IBM MQ V9.1. An M2000 or M2001 appliance can run IBM MQ V8.0 or IBM MQ V9.0, as well as IBM MQ V9.1.

Transfer from an existing single appliance to a new single appliance by using archive files

Follow this procedure to transfer queue managers from an existing IBM MQ Appliance to a new IBM MQ Appliance by using archive files created when backing up each queue manager.

You can back up a queue manager, including log files and data, to an archive file stored in the `mqbackup:///QMGRS` directory on the appliance. You can move the file to another appliance and then restore the queue manager to the new appliance. Follow the instructions in [“Backing up a queue manager”](#) on page 365 and [“Restoring a queue manager”](#) on page 367.

Transfer from an existing single appliance to a new single appliance by using DR commands

Follow this procedure to transfer queue managers from an existing IBM MQ Appliance to a new IBM MQ Appliance using disaster recovery commands.

The following diagram gives an overview of the procedure.

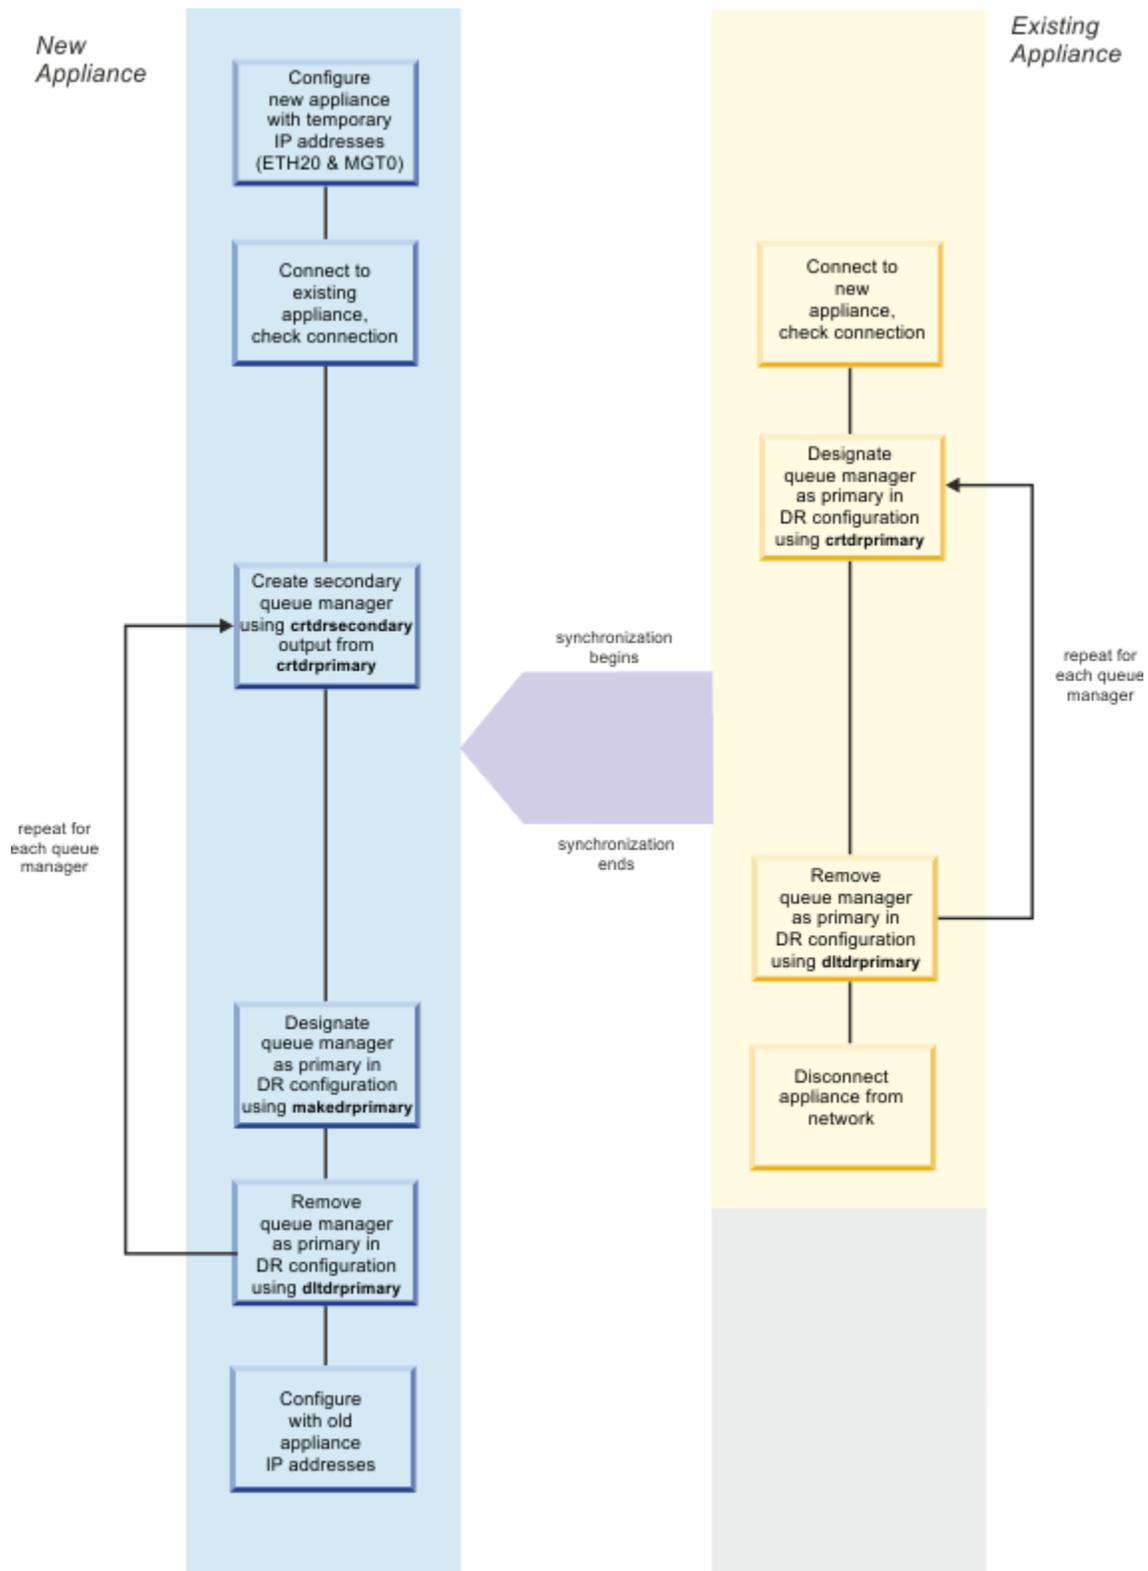


Figure 54. Transferring queue managers on single appliances

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 218](#) 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 218](#) 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_name
```

- c) Enter the following command:

```
crtldrprimary -m queue_manager_name -r new_appliance_name -i new_appliance_IP  
-p port
```

On successful completion, the command outputs a **crtldrsecondary** command.

- d) Restart the queue manager:

```
strmqm queue_manager_name
```

2. On the new appliance, complete the following steps:
 - a) Enter MQ administration mode by typing `mqcli` on the command line.
 - b) Enter the **crtldrsecondary** 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_name
```

Related reference

[“crtldrprimary” on page 680](#)

Augments an existing queue manager to become the primary queue manager in a disaster recovery configuration.

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_name
```

- 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_name
```

- 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_name
```

- c) Stop the queue manager:

```
endmqm queue_manager_name
```

- d) Enter the following command to remove the queue manager from the temporary disaster recovery configuration:

```
dltdrprimary -m queue_manager_name
```

- 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 218 for help with reconfiguring your new appliance.

Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“dltdrprimary” on page 682](#)

Remove a queue manager currently in the primary role from DR control.

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.

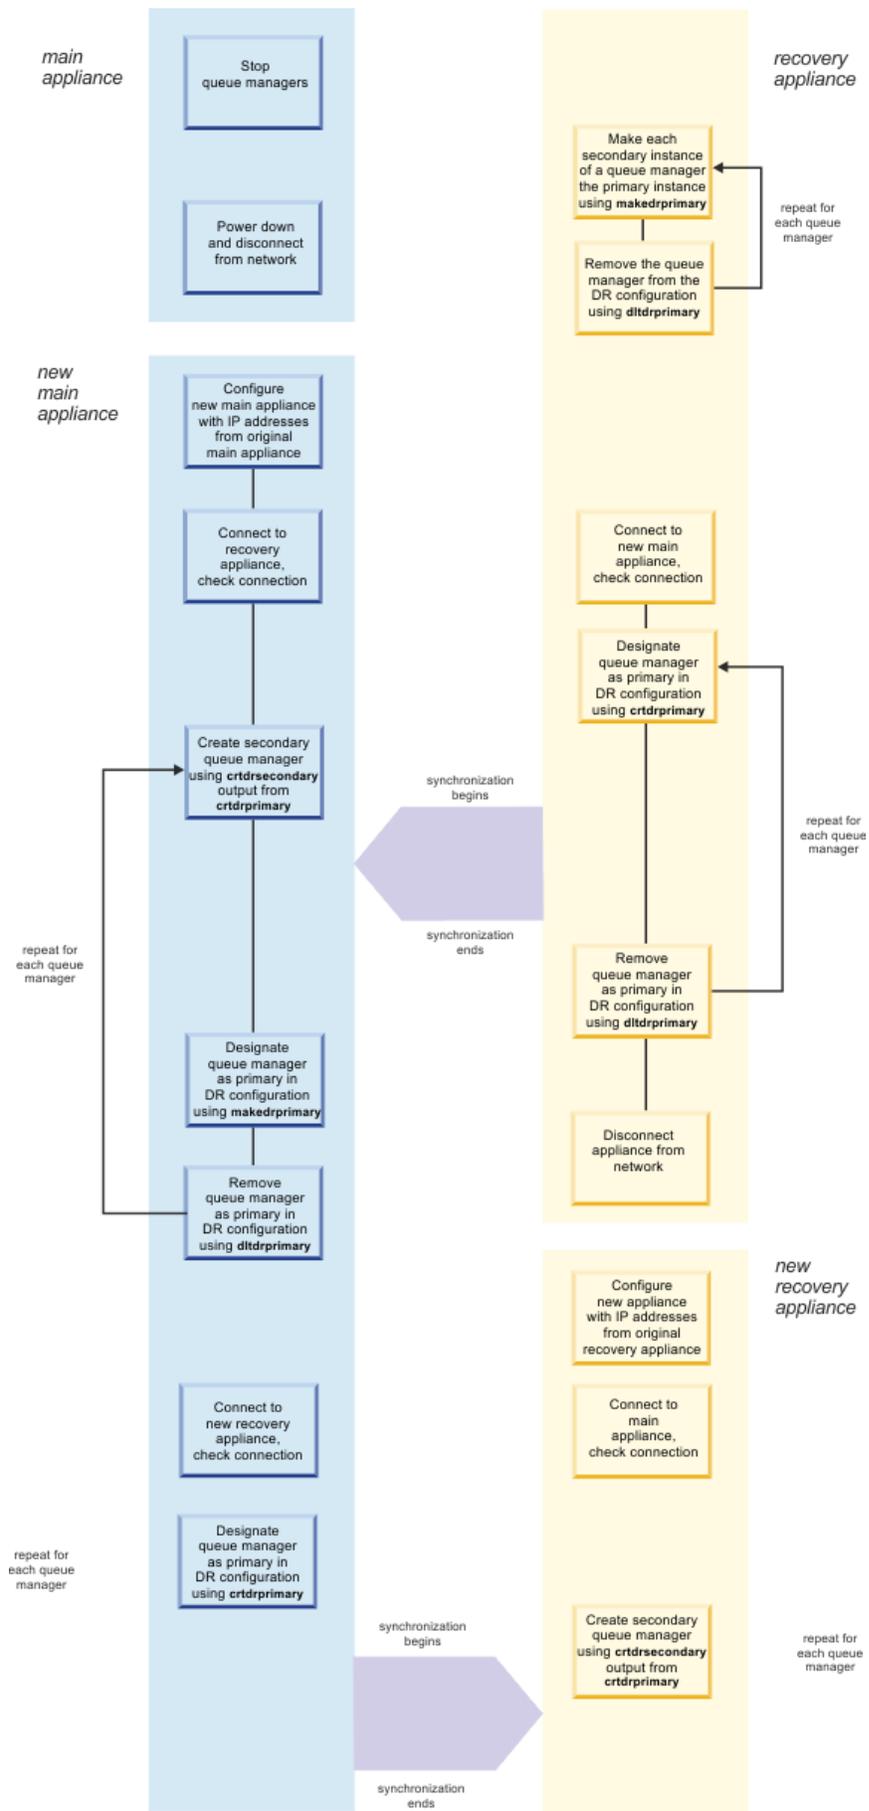


Figure 55. Transferring from an existing appliance in a disaster recovery configuration

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_name
```

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_name
```

- c) Start the queue manager:

```
strmqm queue_manager_name
```

- d) Remove the queue manager from the disaster recovery configuration:

```
dltdrprimary -m queue_manager_name
```

- 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.

Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“dltdrprimary” on page 682](#)

Remove a queue manager currently in the primary role from DR control.

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 218](#) 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_name
```

- c) Enter the following command:

```
crtdrprimary -m queue_manager_name -r new_appliance_name -i new_appliance_IP  
-p port
```

On successful completion, the command outputs a **crtdrsecondary** command.

d) Restart the queue manager:

```
stimqm queue_manager_name
```

2. On the new main appliance, complete the following steps:

a) Enter MQ administration mode by typing `mqcli` on the command line.

b) Enter the **`crtldrsecondary`** 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_name
```

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_name
```

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_name
```

6. Stop each queue manager:

```
endmqm queue_manager_name
```

7. Enter the following command for each queue manager to remove it from the disaster recovery configuration:

```
dltdrprimary -m queue_manager_name
```

Related reference

[“crtldrprimary” on page 680](#)

Augments an existing queue manager to become the primary queue manager in a disaster recovery configuration.

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“dltdrprimary” on page 682](#)

Remove a queue manager currently in the primary role from DR control.

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 218 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_name -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_name
```

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_name
```

12. Repeat steps 10 and 11 for each queue manager.

Related reference

[“crtdrprimary”](#) on page 680

Augments an existing queue manager to become the primary queue manager in a disaster recovery configuration.

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.

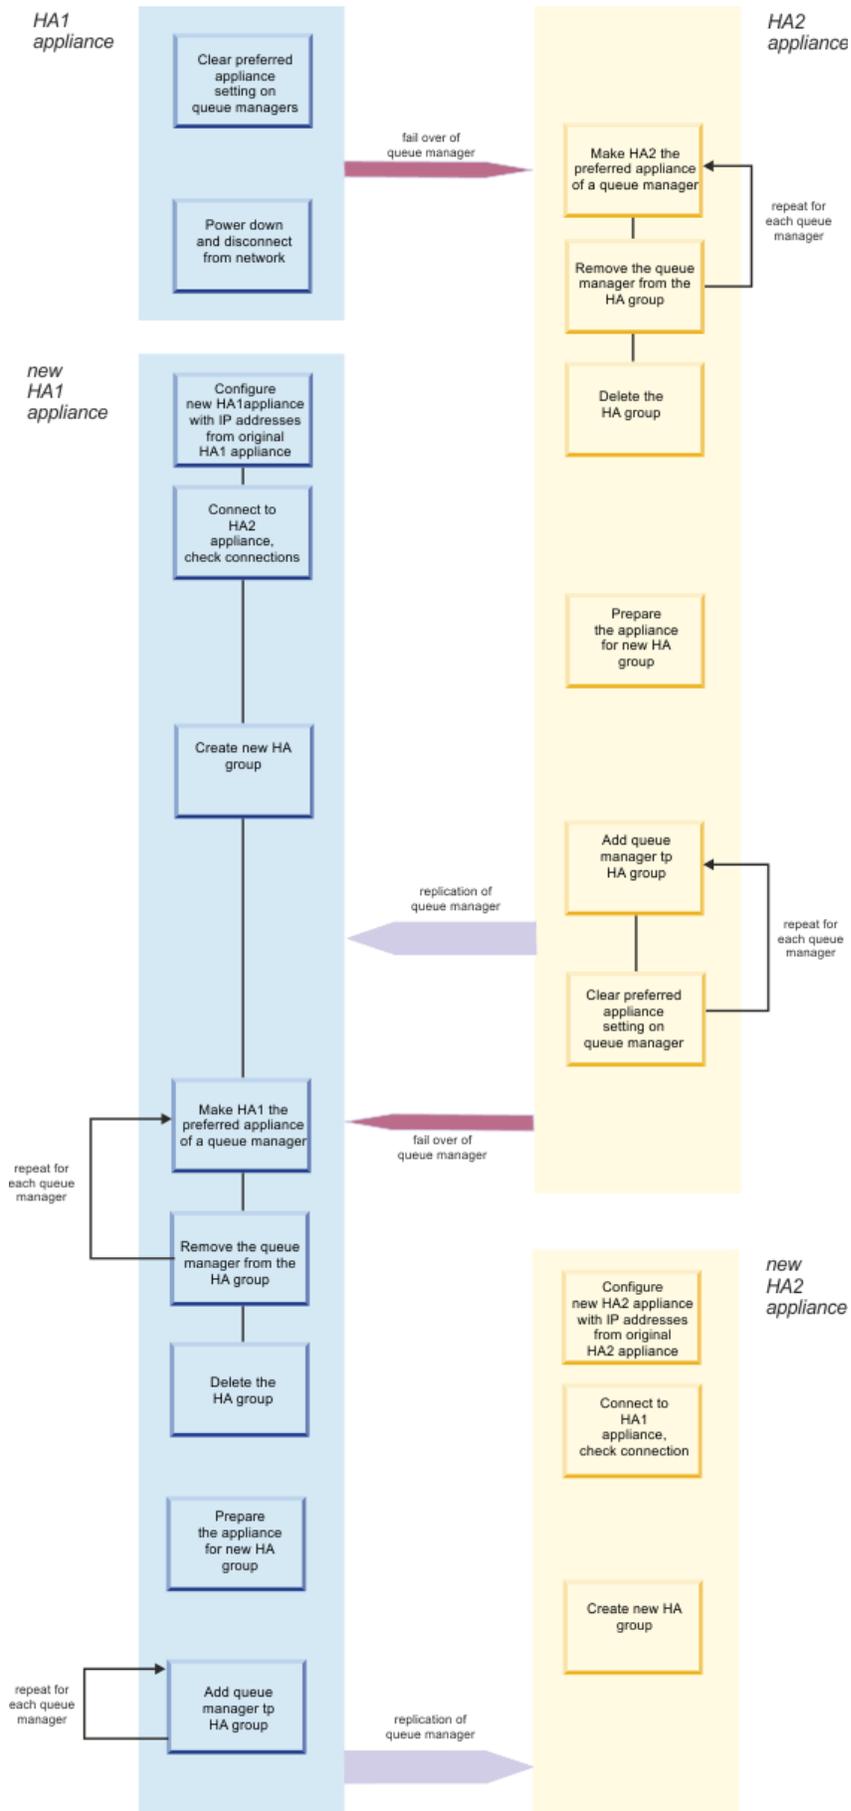


Figure 56. Transferring from an existing high availability group

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_name
```

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_name
```

c) If the queue manager is running, stop it:

```
endmqm queue_manager_name
```

d) Remove the queue manager from the HA group:

```
sethagrp -e queue_manager_name
```

e) Repeat these steps for each queue manager.

f) Delete the HA group:

```
dlthagrp
```

3. Power down appliance HA1 and disconnect all the cables.

Related reference

[“sethapreferred \(set HA preferred appliance\)” on page 679](#)

Sets a preferred appliance in the high availability (HA) group for a queue manager to run on.

[“clearhapreferred \(clear HA preferred appliance\)” on page 679](#)

Clears the preferred appliance for a queue manager in a high availability (HA) group.

[“sethagrp \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

“dlthagrp (delete HA group)” on page 672

Deletes a high availability (HA) group.

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 218 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_name
```

Adding the queue manager to the HA group copies it to the HA1 appliance and restarts it running on the HA2 appliance.

Related reference

[“prepareha \(prepare appliance for HA group\)” on page 674](#)

Prepares an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.

[“crthagrp \(create HA group\)” on page 669](#)

Creates a high availability (HA) group of appliances.

[“sethagrp \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

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_name
```

- c) Clear the preferred appliance setting from the queue manager:

```
clearhappreferred queue_manager_name
```

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:

```
sethappreferred queue_manager_name
```

c) If the queue manager is running, stop it:

```
endmqm queue_manager_name
```

d) Remove the queue manager from the HA group:

```
sethagrp -e queue_manager_name
```

- e) Repeat these steps for each queue manager.
- f) Delete the HA group:

```
dlthagrp
```

3. Power down appliance HA2 and disconnect all the cables.

Related reference

[“sethappreferred \(set HA preferred appliance\)” on page 679](#)

Sets a preferred appliance in the high availability (HA) group for a queue manager to run on.

[“clearhappreferred \(clear HA preferred appliance\)” on page 679](#)

Clears the preferred appliance for a queue manager in a high availability (HA) group.

[“sethagrp \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“dlthagrp \(delete HA group\)” on page 672](#)

Deletes a high availability (HA) group.

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 218](#) 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:

```
crthagip -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:

```
sethagip -i queue_manager_name
```

Adding the queue manager to the HA group copies it to the HA2 appliance and restarts it running on the HA1 appliance.

Related reference

[“prepareha \(prepare appliance for HA group\)” on page 674](#)

Prepares an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.

“crthagrp (create HA group)” on page 669

Creates a high availability (HA) group of appliances.

“sethagr (set HA group)” on page 675

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

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”](#) on page 207.
- Enable SSH and establish a remote connection through the SSH service. For details, see [“Configuring the SSH service”](#) on page 207.

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.

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.

You can also fine tune the ciphers that are used by the SSH service, and the order that they are used in.

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.

About this task

You configure the SSH service, and can then optionally go on to configure the ciphers that the SSH service uses. By default the service uses 16 ciphers in a recommended order. You can, if required, disable or re-enable certain ciphers and reorder them.

Procedure

1. Start the IBM MQ Appliance web UI, and click the network icon .
2. Select **Management > SSH Service**.
3. Select **Enable administrative state**.
4. In the **Local address** field, enter the local address the appliance monitors for incoming SSH requests.
5. In the **Port number** field, change the port on which the appliance monitors for incoming SSH requests.
6. Click **Edit** next to the **Access control list** field to modify the SSH ACL.
7. Click **Apply** to save the changes to the running configuration.
8. Click **Save Configuration** to save the changes to the persisted configuration.
9. If you want to work with the SSH profile ciphers, click the object icon , then select **Crypto configuration > SSH Server profile**.
10. Ensure that **Enable administrative state** is enabled.
11. Select or deselect and reorder the ciphers as required.
12. Click **Apply** to save the changes to the running configuration.
13. 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.

About this task

You configure the SSH service, and can then optionally go on to configure the ciphers that the SSH service uses. By default the service uses 16 ciphers in a recommended order. You can, if required, disable or re-enable certain ciphers and reorder them.

Procedure

1. Enter the configuration mode by entering the following command:

```
config
```

2. 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.

3. Exit the configuration mode by entering the following command:

```
exit
```

4. To work with SSH ciphers, enter the crypto SSH server mode by entering the following commands:

```
crypto  
sshserverprofile
```

5. Enter the following command to change the enabled ciphers and the order of preference that they are used in:

```
ciphers cipher_string
```

Where *cipher_string* lists the enabled ciphers in the required preference order. See [“ciphers” on page 959](#) for the names of ciphers that you can specify in the string.

6. Exit the configuration mode by entering the following command:

```
exit
```

7. Exit the crypto mode by entering the following command:

```
exit
```

Related reference

[“ssh” on page 816](#)

This command enables SSH on appliance interfaces.

[“ciphers” on page 959](#)

This command specifies the cipher suites that the appliance accepts for SSH encryption.

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 210](#)).

Procedure

- To upload certificates to your appliance:
 - a) Ensure that you have the following items:
 - A private key to access the appliance certificate.
 - The appliance certificate.
 - Client certificates (optional).
 - b) Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#).
 - c) Log in as a user in the administrators group.
 - d) Type the following command to enter configuration mode:

```
config
```

e) 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 489](#).

- To create definition objects for the appliance certificate and key:
 - a) From configuration mode, type `crypto` to enter crypto configuration mode.
 - b) 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
```

c) Create a crypto certificate definition for the appliance:

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

For example:

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

d) 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):
 - a) 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
```

b) 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:
 - a) 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:
 - a) Type `exit` to leave crypto configuration mode.
 - b) Type `write mem` to save your configuration changes.
- To associate the SSL server profile with the web UI:
 - a) From configuration mode, type `web-mgmt` to enter web management configuration mode.
 - b) Enter the following command:

```
ssl-server SSL_Svr_Profile_name
```

For example:

```
ssl-server myappliance1
```

- To save your web management configuration:
 - a) Type `exit` to leave web-mgmt configuration mode.
 - b) Type `write mem` to save your configuration changes.
 - c) Type `exit` again to leave configuration mode.

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

[“idcred” on page 733](#)

This command creates identification credentials.

[“valcred” on page 749](#)

This command enters Validation Credentials mode.

[“key” on page 734](#)

This command creates an alias for a private key.

[“ssl-server” on page 748](#)

This command enters SSL Server Profile mode.

Customizing the user interfaces

You can customize certain features of the CLI or web UI on your IBM MQ Appliance.

You can customize the following aspects of the user interfaces:

- The CLI prompt. For example, you could include an identifier for the appliance.
- Pre-login, post-login, and system messages for the CLI.
- Pre-login, post-login, and system messages for the web UI.

You customize the user interfaces using an XML file. You create the file that defines your customizations and then copy it to the appliance.

You can use any text editor to create the XML file. You must cut and paste the markup, and then specify the content of each customized message within the markup.

After the XML file is complete, you can, if you want, validate the conformance of the file against its schema. The schema is available on the appliance under `store:///schemas/dp-user-interface.xsd`.

Copy the file to the appliance store URI, and use the **custom-ui-file** command to implement your customizations, for example:

```
mqa# config
mqa (config)# system
mqa (config system)# custom-ui-file store:///CustomUI.xml
```

Related reference

[“custom-ui-file” on page 977](#)

This command specifies the location of the custom user interface file.

[“copy” on page 788](#)

This command copies a file to or from the appliance.

Supported markup for the user interface customization file

The user interface customization file is an XML file.

You can copy and paste elements from the supplied template file to create the file, (see [“Template of the custom user interface file” on page 216](#)). The schema for the XML file supports the following case-sensitive elements:

<User-Interface>

The <User-Interface> element is the root element of the XML file and defines the required namespace statements. The XML file must contain this element copied and pasted from the template without modification.

<CustomPrompt>

The <CustomPrompt> element indicates whether to extend the CLI prompt with the system name. To enable this aspect, add an element of the form:

```
<CustomPrompt>%s</CustomPrompt>
```

The system name is the only customization available for the CLI prompt. The system name can be set using the **name** command, see [“name” on page 979](#).

<MarkupBanner>

The <MarkupBanner> element identifies the messages to display in the web UI. The file can contain up to four <MarkupBanner> elements, based on a combination of the `type` attribute and `location` attribute.

type="message-type"

The `type` attribute identifies the type of message. This attribute supports the following keywords.

pre-login

Displays the message before users log in to the web UI. You can define one pre-login message.

post-login

Displays the message in a pop-up window immediately after users log in to the web UI. You can define one post-login message.

system-banner

Displays the message on each web UI screen. You can define two appliance messages based on the keyword of the `location` attribute. Use the `location` attribute to define where on the web UI to display the message.

location="location"

The `location` attribute indicates the location on the web UI to display the message. This attribute is relevant only when used with `type="system-banner"`. The `location` attribute supports the following keywords. The default value is `both`.

header

Displays the message at the top. You can define one message with this keyword. You cannot define a message with this keyword and another with the `both` keyword.

footer

Displays the message at the bottom. You can define one message with this keyword. You cannot define a message with this keyword and another with the `both` keyword.

both

Displays the message at the top and the bottom. You can define one message with this keyword. You cannot define a message with this keyword and another with the `header` keyword or with the `footer` keyword.

foreground-color="color"

The `foreground-color` attribute identifies the color of the text in the web UI message. This attribute supports the following keywords. The default value is `none`, which display the text in black.

- none
- blue
- green
- orange
- red
- yellow

background-color="color"

The `background-color` attribute identifies the color of the background in the web UI message. This attribute supports the following keywords. The default value is `none`, which removes any color from the message background.

- none
- blue
- green
- orange
- red

- yellow

For web UI messages, the contents of the <MarkupBanner> element can include the following standard HTML tags.

<p>

Defines individual paragraphs.

Defines text to display in italics.

Defines text to display in bold.

<tt>

Defines text to display in monospace.

<TextBanner>

The <TextBanner> element identifies the messages to display to users in CLI sessions. The file can contain up to three <TextBanner> elements, one for each keyword that is associated with the type attribute.

type="message-type"

The type attribute identifies the type of message. This attribute supports the following keywords.

pre-login

Displays the message before users log in from the CLI.

post-login

Displays the message immediately after users log in from the CLI.

system-banner

Displays the message immediately after the completion of each command invocation from the CLI.

For CLI messages, the content of the <TextBanner> element cannot include other HTML or XML elements.

Template of the custom user interface file

You can use the provided template as a base for producing a custom user interface file.

The following template is an XML file to help you create the custom user interface file for your IBM MQ Appliance. This template conforms to the schema (store:///schemas/dp-user-interface.xsd).

```
<User-Interface
  xmlns="http://www.datapower.com/schemas/user-interface/1.0">

  <!-- Markup for the CLI prompt -->
  <CustomPrompt>%s</CustomPrompt>

  <!-- Markup for custom messages for the GUI -->
  <MarkupBanner type="pre-login" foreground-color="red" background-color="blue">
    GUI pre-login message
  </MarkupBanner>
  <MarkupBanner type="post-login" foreground-color="blue" background-color="yellow">
    GUI post-login pop up message
  </MarkupBanner>
  <MarkupBanner type="system-banner" location="header" foreground-color="green"
    background-color="red">
    GUI system message - header
  </MarkupBanner>
  <MarkupBanner type="system-banner" location="footer" foreground-color="blue"
    background-color="yellow">
    GUI system message - footer
  </MarkupBanner>

  <!-- If the following markup was outside of comments, the file would not
    conform to the schema. Cannot define multiple system messages as the
    header or footer.
  <MarkupBanner type="system-banner">
    GUI system message - header and footer
  </MarkupBanner>
  -->
```

```
<!-- Markup for custom messages for the CLI -->
<TextBanner type="pre-login">
  CLIE pre-login message
</TextBanner>
<TextBanner type="post-login">
  CLI post-login message
</TextBanner>
<TextBanner type="system-banner">
  CLI system message
</TextBanner>
</User-Interface>
```

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. 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.
2. 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 269](#) and [“Changing IP addresses in disaster recovery configurations” on page 289](#).
3. 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 Network section 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 269](#) and [“Changing IP addresses in disaster recovery configurations” on page 289](#).

Procedure



1. Start the IBM MQ Appliance web UI, and click the network icon.
2. Select **Interface > 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 autoconfiguration 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.

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 269](#) and [“Changing IP addresses in disaster recovery configurations” on page 289](#).

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `conf` 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 708	Sets the administrative state for the configuration.
“flow-control” on page 770	Sets the flow control mode of the Ethernet interface.
“ip-address” on page 772	Assigns the primary network address for the Ethernet interface.
“ip-config-mode” on page 773	Identifies the configuration mode for the Ethernet interface.
“ip-route” on page 774	Manages static routes in the routing table for the Ethernet interface.
“ip-secondary-address” on page 775	Manages secondary network addresses for the Ethernet interface.
“ipv4-default-gateway” on page 776	Designates the default IPv4 gateway for the Ethernet interface.
“ipv6-dadtransmits” on page 776	Sets the number of IPv6 duplication address detection attempts for the Ethernet interface.
“ipv6-default-gateway” on page 777	Designates the default IPv6 gateway for the Ethernet interface.
“ipv6-nd-retransmit-timer” on page 777	Sets the interval between IPv6 neighbor discovery attempts for the Ethernet interface.
“link-aggregation-mode” on page 778	Indicates whether the Ethernet interface is part of an aggregate interface.
“mac-address” on page 778	Changes the MAC address for the Ethernet interface.
“mode” on page 779	Sets the interface speed and direction.
“mtu” on page 779	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 771	Manages the temporary disabling of hardware offload.
“packet-capture” on page 780	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.

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 Network section 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 the network icon .
2. Select **Interface > 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 autoconfiguration 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 define advanced options that modify the way that the VLAN interface works within the network.
11. Click **Apply** to save the changes to the running 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 207](#). Log in as an administrative user.
2. Type `conf` 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 708	Sets the administrative state for the configuration.
“ethernet-interface” on page 987	Sets the Ethernet interface to provide connectivity.
“identifier” on page 987	Sets the VLAN identifier.
“ip-address” on page 988	Assigns the primary network address for the VLAN interface.
“ip-config-mode” on page 989	Identifies the configuration mode for the VLAN interface.
“ip-route” on page 990	Manages static routes in the routing table for the VLAN interface.
“ip-secondary-address” on page 991	Manages secondary network addresses for the VLAN interface.

Command	Description
“ipv4-default-gateway” on page 991	Designates the default IPv4 gateway for the VLAN interface.
“ipv6-dadtransmits” on page 992	Sets the number of IPv6 duplication address detection attempts for the VLAN interface.
“ipv6-default-gateway” on page 992	Designates the default IPv6 gateway for the VLAN interface.
“ipv6-nd-retransmit-timer” on page 993	Sets the interval between IPv6 neighbor discovery attempts for the VLAN interface.
“link-aggregation-interface” on page 993l	Indicates whether the VLAN interface is part of an aggregate interface.
“mtu” on page 994	Sets the maximum transmission unit of the VLAN interface.
“outbound-priority” on page 994	Sets the priority value in outbound packets.
“over” on page 995	Sets the parent interface type.

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

Command	Description
“packet-capture” on page 996	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.

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 Network section 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 the network icon .
2. Select **Interface > 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.

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 207. 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 708	Sets the administrative state for the configuration.
“ip-address” on page 838	Assigns the primary network address for the link aggregation interface.
“ip-config-mode” on page 839	Identifies the configuration mode for the link aggregation interface.
“ip-route” on page 840	Manages static routes in the routing table for the link aggregation interface.
“ip-secondary-address” on page 841	Manages secondary network addresses for the link aggregation interface.

Command	Description
“ipv4-default-gateway” on page 842	Designates the default IPv4 gateway for the link aggregation interface.
“ipv6-dadtransmits” on page 842	Sets the number of IPv6 duplication address detection attempts for the link aggregation interface.
“ipv6-default-gateway” on page 843	Designates the default IPv6 gateway for the link aggregation interface.
“ipv6-nd-retransmit-timer” on page 843	Sets the interval between IPv6 neighbor discovery attempts for the link aggregation interface.
“lACP-hash” on page 843	Sets which hash function the LACP aggregation uses to determine the interface for outbound packets.
“lACP-select” on page 844	Sets the algorithm for the LACP selection policy.
“link” on page 845	Specifies which Ethernet interfaces are part of the aggregate interface.
“type” on page 847	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 846	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.



CAUTION: Enabling IPMI over LAN exposes security vulnerability [CVE-2013-4786](#), therefore IPMI over LAN must only be enabled in a secure environment.

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 `mg0` 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 207](#). 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 708	Sets the administrative state for the configuration.
“allowed-user” on page 831	Specifies the users allowed to use the channel.
“ip address” on page 832	Sets IP addresses with subnet mask for the IPMI LAN channel.
“ip default-gateway” on page 833	Sets the default gateway for the IPMI LAN channel.
“maximum-channel-privilege-level” on page 833	Sets the maximum privilege level for users.
“sol-enabled” on page 834	Indicates whether to support serial over LAN.
“sol-required-user-privilege-level” on page 834	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 207](#). 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 settings 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 the network icon .
2. Select **Interface > 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.

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 207. 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 708	Sets the administrative state for the configuration.

Command	Description
“ip-preference” on page 765	Sets the preferred IP version when the DNS provider publishes both versions of addresses.
“load-balance” on page 766	Sets the load-balancing algorithm that the appliance uses to resolve host names.
“name-server” on page 766	Manages local DNS providers.
“retries” on page 767	Sets the number of times that the appliance attempts a failed query.
“search-domain” on page 767	Manages domain-suffixes in the search table for non-qualified domain names.
“static-host” on page 768	Manages host-address maps.
“timeout” on page 769	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 the network icon .
2. Select **Interface > DNS Service**.
3. Click **Actions > Flush DNS Cache**.

SNMP Settings

The IBM MQ Appliance supports SNMP versions 1, 2c, and 3.

SNMP versions 1 and 2c support a community-based security model. SNMP version 3 has a security model that uses SHA or MD5 based authentication and AES or DES encryption. The two security models use different terminology: SNMP version 2 and earlier uses the term "traps" for messages that inform about events, while SNMP version 3 calls these messages "notifications".

You can configure the SNMP settings for the appliance by using the IBM MQ Appliance web UI or the command line interface. The web UI settings are divided into the following sections:

Main

The main SNMP settings specify local connection and security settings. You specify the local IP address and port that is listened on for incoming SNMP requests. You can specify that all interfaces on the appliance are listened on, if required. If you are using SNMP v3, you can specify the SNMP users, security levels, and access levels.

Enterprise MIBs

You can view the three MIB files that specify the interaction of the appliance with SNMP managers. You need to download the MIB files to use them with an SNMP monitoring application. All of the objects that are exposed through the IBM MQ Appliance MIBs are read only.

Trap Event Subscriptions

Enables the generation of event traps/notifications for certain appliance conditions. Currently, you can enable or disable the generation of traps/notifications in response to a default set of appliance events. Note that events for IBM MQ operations are not currently supported.

SNMPv1/v2c Communities

If you are using SNMP version 1 or version 2c, you use this section to specify one or more SNMP communities.

Trap and Notification Targets

If you have enabled the default event traps/notifications, you use this tab to specify details of the SNMP manager (or managers) that the traps/notifications are sent to.

Specifying main SNMP settings by using the web UI

Enable SNMP and specify main connection and security details.

About this task

You can use the IBM MQ Appliance web UI to configure SNMP. You use the **Main** section of the interface to specify basic details and to enable the service.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > SNMP Settings**.
3. Ensure that **Enable administrative state** is selected (this step enables SNMP).
4. Specify the IP address of the local interface on which the SNMP service listens for SNMP requests in the **Local IP Address** field. Specify 0.0.0.0 (which is the default value) to listen on all appliance interfaces.
5. Specify the port that is listened on in the **Local Port** field. The port for SNMP is 161 by default.
6. If you are configuring SNMP v3, complete the following steps:
 - a) If required by the chosen security level, add one or more local users who have SNMPv3 credentials:
 - If you have already configured a local user and specified SNMPv3 credentials, click **Add** and select the user from the list. Repeat this step to add additional users.
 - If you have already configured a local user, but not yet specified their specified SNMPv3 credentials, click **Add** and select the user from the list, then click the pencil icon  to edit the user definition. Add the required credentials and click **Apply**.
 - Otherwise, click the plus icon  to open the User Account dialog. Create a new user and specify SNMPv3 credentials. Click **Apply**.
 - See [“Configuring local users by using the web UI” on page 471](#) .
 - b) Select the SNMPv3 security level. Choose one of the following options:
 - Authentication, Privacy**
The SNMP connection requires authentication of users and encryption of data. This setting is the default.
 - Authentication, No Privacy**
The SNMP connection requires authentication of users but not the encryption of data.
 - No Authentication, No privacy**
The SNMP connection requires neither authentication of users nor encryption of data.
 - c) Select the SNMPv3 access level. You should select the **read-only** option, which allows an SNMP manager to request get, get-next, and get-bulk operations. The **read-only** option is the default. (All of the objects that are exposed through the IBM MQ Appliance MIBs are read only.)

What to do next

If you want to view the MIBs that control how SNMP managers work with the appliance, open the **Enterprise MIBs** section of the dialog, or open the **Trap Event Subscriptions** section to configure default

traps/notifications. To specify communities for SNMPv1 or SNMP v2c, open the **SNMPv1/v2c Communities** section. Otherwise, if you have completed your SNMP configuration, click **Apply**.

Viewing MIBs by using the web UI

You can view MIBs, but you cannot change them.

About this task

The interaction of an SNMP manager with the appliance is controlled by MIB files, which define the objects that the SNMP manager can collect data about. There are three MIBs that control this interaction: configuration, status, and notifications.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > SNMP Settings**.
3. Open the **Enterprise MIBs** section of the dialog.
4. Click the link for the MIB that you want to view. The MIB opens in a separate browser window.

What to do next

Open the **Trap Events** section to configure the default events that are generated by the appliance.

Configuring default trap events by using the web UI

The appliance has a set of default events that can be reported if you enable this feature.

About this task

The default events report on the functioning on the appliance itself rather than the operation of IBM MQ. You can enable or disable the generation of the default events as a whole, and specify the level at which an event is generated for the specified conditions.

The default events are specified by code. The codes correspond to the following conditions:

0x00030002

Out of memory

0x00230003

Unable to allocate execution resources

0x00330002

Memory full

0x00b30014

Duplicate IP address

0x00e30001

NTP - Cannot Resolve Server Name

0x00e40008

NTP Timeout Error

0x00f30008

File is expired (refers to certificate file)

0x01530001

Time zone config mismatch

0x01a2000e

Installed battery is nearing end of life

0x01a40001

Throttling connections due to low memory

0x01a40005

Throttling connections due to low temporary file space

0x01a40008

Throttling connections due to low number of free ports

0x01b10009

Uncertified HSM firmware detected (not used)

0x01b20002

HSM is uninitialized (not used)

0x01b20004

HSM PED login failed (not used)

0x01b20008

HSM password login failed (not used)

0x02220001

Power supply failure

0x02220003

Internal cooling fan has stopped

0x02240002

Internal cooling fan has slowed

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > SNMP Settings**.
3. Open the **Trap Event Subscriptions** section of the dialog.
4. Select **Enable Default Event Subscriptions** to enable the generation of notification of trap events.
5. Select a minimum trap event priority from the **Minimum Priority** list. The priorities are hierarchical. The lowest is listed last. Set to the minimum level of event that will cause a trap/notification to be generated.

emergency

An emergency level message. The system is unusable.

alert

An alert level message. Immediate action must be taken.

critical

A critical message. Immediate action should be taken.

error

An error message. Processing might continue, but action should be taken.

warning

A warning message. Processing should continue, but action should be taken.

notice

A notice message. Processing continues, but action might need to be taken.

information

An information message. No action required.

debug

A debug message for processing information to help during troubleshooting.

What to do next

If you are using SNMP v1 or v2c, open the **SNMPv1/v2c Communities** section to specify community details. Otherwise, open the **Trap and Notifications Targets** section to specify the destination for the default trap notifications, if you enabled them.

Configuring communities by using the web UI

If you are using SNMP version 1 or version 2c, you must specify one or more communities.

About this task

The v1 and v2c versions of SNMP rely on a password phrase that is known as a "community". The name of the community accompanies SNMP requests, and is used to determine whether the request can be fulfilled or not.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > SNMP Settings**.
3. Open the **SNMPv1/v2c Communities** section of the dialog.
4. Click **Add** to add an SNMP community.
5. Specify the name of the community in the **Community** field.
6. Select the type of access that SNMP managers using this community name have to the appliance from the **Mode** list.
7. Specify the IP address of the SNMP manager for the appliance in the **Remote Host Address** field. Set the address to 0.0.0.0/0 to allow access to all SNMP managers that use the community name.
8. If required, click **Add** to add more communities.

What to do next

If you have specified that the appliance can generate trap event notifications, open the **Trap and Notifications Targets** section to specify where notifications are sent. Otherwise, if you have completed your SNMP configuration, click **Apply**.

Configuring trap and notification targets by using the web UI

If you enable default trap events, you should also define a target to send them to.

About this task

The information that you supply when you configure a target depends on the version of SNMP that you are using. SNMPv1 and v2c require different setup information to SNMP v3.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > SNMP Settings**.
3. Open the **Trap and Notification Targets** section of the dialog.
4. Specify the IP address of the target recipient of traps/notifications in the **Remote Host Address** field.
5. Specify the port to use when sending traps/notifications in the **Remote Port** field. SNMP uses port 162 by default.
6. Select the version of SNMP you are using from the **Version** list.
 - If you select version 1 or 2c, specify the name of the community to use when sending trap events in the **Community** field.

- If you select version 3, specify the name of the SNMP user used to send event notifications in the **Security Name** field, and select the security level from the **Security Level** list to specify whether the user is authenticated and data is encrypted.

What to do next

If you have completed your SNMP configuration, click **Apply**.

Configuring SNMPv3 settings by using the command line

Enable SNMP and specify main connection and security and user details.

About this task

You can use the command line to configure connection details for SNMPv3.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type `snmp` to enter SNMP configuration mode.
4. Enable SNMP by entering the following command:

```
admin-state enabled
```

5. Specify the IP address of the local interface on which the SNMP service listens for SNMP requests by entering the following command:

```
ip-address local_IP_address
```

Where *local_IP_address* is a local IP address. Specify 0.0.0.0 to listen on all appliance interfaces.

6. Specify the port that is listened on by entering the following command:

```
port port_number
```

Where *port_number* is the port listened on. The port is set to 161 by default.

7. Specify the security level by entering the following command:

```
security-level level
```

Where *level* is one of the following values:

noAuthNoPriv

The SNMP connection requires neither authentication of users nor encryption of data.

authNoPriv

The SNMP connection requires authentication of users but not the encryption of data.

authPriv

The SNMP connection requires authentication of users and encryption of data.

If you select either of the levels that specify user authentication, you must define a local user for SNMP authentication, together with SNMP credentials. See [“Configuring local users by using the command line” on page 472](#).

8. Specify the user ID of the local user that is used for authentication:

```
user userName
```

9. Specify the access level of read-only by entering the following command:

```
access-level read-only
```

What to do next

Specify whether default trap events are enabled.

Related reference

[“ip-address” on page 953](#)

This command specifies the IP address that is listened on for SNMP requests.

[“port” on page 954](#)

This command identifies the appliance UDP port that is monitored by the SNMP agent for SNMP requests.

[“security-level” on page 954](#)

This command specifies the access that an SNMPv3 manager has to the appliance.

[“user” on page 958](#)

This command specifies the user ID used by a SNMPv3 manager to connect to the appliance.

[“access-level” on page 952](#)

This command specifies the level of access that an SNMPv3 manager has to the appliance MIBs.

Configuring SNMPv1 or v2c settings by using the command line

Enable SNMP and specify main connection and community details.

About this task

The v1 and v2c versions of SNMP rely on a password phrase that is known as a "community". The name of the community accompanies SNMP requests, and is used to determine whether the request can be fulfilled or not.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `confi`g to enter global configuration mode.
3. Type `snmp` to enter SNMP configuration mode.
4. Enable SNMP by entering the following command:

```
admin-state enabled
```

5. Specify the IP address of the local interface on which the SNMP service listens for SNMP requests by entering the following command:

```
ip-address local_IP_address
```

Where *local_IP_address* is a local IP address. Specify 0.0.0.0 to listen on all appliance interfaces.

6. Specify the port that is listened on by entering the following command:

```
port port_number
```

Where *port_number* is the port listened on. The port is set to 161 by default.

7. Enter the following command to specify a community:

```
community communityName access-level [ip_address]
```

Where:

communityName

Specifies the name of the community.

access-level

Set *access-level* to `read-only` to specify that SNMP managers are restricted to SNMP get operations, which means that these managers can read, but cannot change management information base (MIB) values.

ip_address

Optionally, specify an IP address to restrict access to the SNMP manager in the named community with the specified IP address. By default, any SNMP manager that belongs to the named community can make requests.

What to do next

Specify whether default trap events are enabled.

Related reference

[“ip-address” on page 953](#)

This command specifies the IP address that is listened on for SNMP requests.

[“port” on page 954](#)

This command identifies the appliance UDP port that is monitored by the SNMP agent for SNMP requests.

[“community” on page 953](#)

This command grants and defines access to the specified SNMPv1 or v2c communities.

Configuring default trap events by using the command line

Enable this feature to report default events to an SNMP if they arise.

About this task

You can use the command line to enable or disable the reporting of default trap events and to specify a target for them.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type `snmp` to enter SNMP configuration mode.
4. Enable default trap events by entering the following command:

```
trap-default-subscriptions on
```

Specify `off` to disable the default trap events.

5. Specify the version of SNMP that you are using:

```
version versionNumber
```

Where *versionNumber* can be 1, 2c, or 3.

6. If you are using SNMP v3, enter the following command:

```
trap-target ip_address [port] snmpVersion user security_level
```

Where:

ip_address

Specifies the IP address of the host that receives notifications.

port

Optionally specify the port. The port is 162 by default.

snmpVersion

Specify 3 to indicate that you are using SNMPv3.

user

Specify the user ID that is used to send notifications.

security_level

Specify the security level. The level is one of `noAuthNoPriv`, `authNoPriv`, or `authPriv`.

7. If you are using SNMPv1 or SNMP v2c, enter the following command:

```
trap-target ip_address [port] snmpVersion community
```

Where:

ip_address

Specifies the IP address of the host that receives notifications.

port

Optionally specify the port. The port is 162 by default.

snmpVersion

Specify 1 or 2c to indicate that you are using SNMPv1 or SNMPv2c.

community

Provides a community name, which is essentially a password, to include in the SNMP message header. The default value is public.

What to do next

If you have finished configuring SNMP, type `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Related reference

[“trap-default-subscriptions” on page 956](#)

This command enables or disables the default list of event codes that generate traps.

[“version” on page 958](#)

This command specifies the supported SNMP version.

[“trap-target” on page 957](#)

This command specifies the recipient of SNMP traps issued by the local SNMP agent.

SNMP validation

You can validate and troubleshoot your SNMP settings by configuring a log target.

If your SNMP monitoring software does not receive SNMP notifications from the appliance, you can test SNMP notifications by configuring an SNMP log target. With an SNMP log target, you can subscribe to reproducible events, such as commands that you enter. The notifications are sent to the SNMP monitoring software that you defined in your SNMP settings.

Complete the following tasks to validate your SNMP settings:

1. Create an SNMP log target that subscribes to CLI events.
2. Test that the SNMP log target captures the CLI events.
3. Add event subscriptions to the SNMP settings for messages that you want to capture.

If further troubleshooting is required, you can capture SNMP traffic with the packet capture feature.

[Creating an SNMP log target](#)

Create an SNMP log target to test SNMP notifications.

Before you begin

The SNMP settings must be configured before you can use an SNMP log target, see [“SNMP Settings” on page 227](#).

About this task

To test SNMP notifications, create an SNMP log target that triggers an event for any command that you enter from the CLI. Define the SNMP log target with an event category of `cli` and an event priority of `debug`. Check the log target status for the number of events that this logging target processed.

Procedure

1. Start the IBM MQ Appliance web UI and click the Objects icon .
2. Select **Logging Configuration > Log Target**.
3. Click **New**.
4. Define the basic properties: Name, administrative state, and descriptive summary.
5. From the **Target Type** list, select **SNMP**.
6. Click **Event Subscriptions**.
 - a) Click **Add**.
 - b) From the **Event Category** list, select **cli**.
 - c) From the **Minimum Event Priority** list, select **debug**.
 - d) Click **Apply**.
7. Click **Apply** to save the changes to the running configuration.
8. Click **Save Configuration** or **Save changes** to save the changes to the persisted configuration.
9. Optionally check the log target status:
 - a) Click the status icon .
 - b) Select **Main > Log Target Status**.
 - c) In the logging targets catalog, note the number of processed events.

Results

With the configured SNMP settings and this SNMP log target, each command that you enter triggers a notification to your SNMP monitoring software. Also, each command increments the number of processed events in the logging target status.

What to do next

Test the SNMP log target to generate CLI log events.

Testing the SNMP log target

Test that your SNMP log target captures all commands and CLI log events.

About this task

To verify that your SNMP log target is active and that it captures events, test the capture of all commands. Test whether commands generate events in the CLI log category.

Procedure

- Enter a command from a CLI session.
 - a) Log on to the appliance.
 - b) Type any command. For example, **show memory**.
- Generate a log event from the Web UI.
 - a) Start the IBM MQ Appliance web UI and click the Administration icon .
 - b) Select **Debug > Troubleshooting**.
 - c) Locate the **Generate Log Event** area.
 - a. From the **Log Category** list, select **cli**.
 - b. From the **Log Level** list, select **debug**.

- c. In the **Log Message** field, enter a message.
- d. Click **Generate Log Event**.
- d) Optionally check the log target status:
 - a. Click the status icon .
 - b. Select **Main > Log Target Status**.
 - c. In the logging targets catalog, note the number of processed events.

Results

With these steps, you generated two CLI events with two different methods. Because you have an SNMP log target that subscribes to CLI events, two SNMP traps are sent to your SNMP notification target. Also, each CLI event increments the number of processed events in the logging target status.

What to do next

Add event subscriptions to the SNMP settings for events that you have a requirement to monitor.

Adding event subscriptions

Add event subscriptions to expand the types of SNMP notifications that you want to test.

Before you begin

The SNMP Settings must be configured, and you must have an SNMP log target.

About this task

In addition to trapping events that are generated, you can trap events that occur from normal operations. Review your system log for events that you want to trap. Configure SNMP notifications with message IDs that are in the system log.

Procedure

1. View the system logs and identify messages that represent events that you want to trap (see [“IBM MQ error logs” on page 499](#)).
2. For each message, copy the value of the msgid and associated error message to a text file.
3. In the IBM MQ Appliance web UI, click the Administration icon  and select **Access > SNMP Settings**.
4. Click the **Trap Event Subscriptions** tab.
5. In the **Minimum Priority** field, select **debug** for the maximum output while testing.
6. Add event subscriptions.
 - a) In the **Event Subscriptions** field, enter a message ID that you copied from the system log.
 - b) Click **Add**.
 - c) Optionally, click **Help** to see the description, level, and suggestion for the message that you added. Confirm that the help text describes the same event that is in your system log.
7. If required, repeat the previous step to add other messages to trap.
8. Click **Apply** to save the changes to the running configuration.
9. Click **Save Configuration** or **Save changes** to save the changes to the persisted configuration.

Results

SNMP notifications are sent for the events that match the event subscriptions.

What to do next

Review the available event codes and look for the events that you have a requirement to monitor. The SNMP settings configuration has a predefined default list of event to trap. You can force errors to test some events, such as configuring an NTP server with a host name that cannot be reached on the network.

If the appliance does not correctly generate SNMP notifications, use a packet capture to troubleshoot.

Troubleshooting SNMP notifications

Use the appliance packet capture for troubleshooting SNMP notifications.

The appliance has a packet capture feature on the Troubleshooting pane. Click the Administration icon



and select **Debug > Troubleshooting** to open the **Troubleshooting** pane.

SNMP traffic can be seen in a packet capture. If the appliance fails to properly generate SNMP notifications, start a packet capture, and issue commands to generate SNMP traps. IBM Support can review the packet capture and help you to understand what data is flowing to your SNMP notification target.

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 the network icon .
2. Select **Interface > 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.  In the **Timeout** field, specify the duration to wait for a response from an NTP server.
6. Click **Apply** to save the changes to the running 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 207](#). 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 708	Sets the administrative state for the configuration.
“refresh-interval” on page 763	Sets the interval between clock synchronizations.
“remote-server” on page 763	Identifies an NTP server.
“timeout” on page 764	Sets the duration to wait for a response from an NTP server

5. After you configure the NTP service settings, enter `exit`
6. Enter `write memory` to save the updated configuration.

Related reference

[“ntp-service” on page 807](#)

This command enters NTP service mode.

Configuring the timezone by using the IBM MQ Appliance web UI

You can 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 the administration icon .
2. Select **Device > Time Settings**.
3. Select the timezone from the **Local time zone** list.
4. Click **Apply** to save the changes to the running 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 207](#). 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 708	Sets the administrative state for the configuration.
“custom” on page 751	Specifies the name of a custom time zone.
“daylight-name” on page 751	Specifies the name of the time zone when in Daylight Saving Time. This name is appended to the time.
“daylight-offset” on page 751	Sets the offset, in hours, for Daylight Saving Time.
“daylight-start-day” on page 752	Specifies the day of the week when Daylight Saving Time starts.
“daylight-start-hours” on page 752	Specifies the hour in the day when Daylight Saving Time starts.
“daylight-start-minutes” on page 753	Specifies the minute in the hour when Daylight Saving Time starts.
“daylight-start-month” on page 754	Specifies the month of the year when Daylight Saving Time starts.
“daylight-start-week” on page 755	Specifies the instance of the day in the month when Daylight Saving Time starts.
“daylight-stop-day” on page 756	Specifies the day of the week when Daylight Saving Time stops.
“daylight-stop-hours” on page 756	Specifies the hour in the day when Daylight Saving Time ends.
“daylight-stop-minutes” on page 757	Specifies the minute in the hour when Daylight Saving Time ends.
“daylight-stop-month” on page 758	Specifies the month of the year when Daylight Saving Time ends.
“daylight-stop-week” on page 759	Specifies the instance of the day in the month when Daylight Saving Time ends.
“direction” on page 760	Specifies the direction, relative to Coordinated Universal Time, of the time zone.
“name” on page 760	Sets the name of the time zone. This name is appended to the time.
“offset-hours” on page 762	Specifies the offset in hours, relative to Coordinated Universal Time, of the time zone.
“offset-minutes” on page 762	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 the administration icon.
2. Select **Device > 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 **Device > System Settings**.
6. Select the locale from the **System locale** list.
7. Click **Apply** to save the changes to the running 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 207](#). 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 `exit` to leave config mode.
 5. Type the following command to enter system settings mode:

```
system
```

6. Type the following command to specify the locale:

```
locale language_code
```

7. After you configure the locale settings, enter `exit`
8. Enter `write memory` to save the updated configuration.

Related reference

[“language” on page 801](#)

This command enters the Language mode to set the administrative state of a language.

[“locale” on page 978](#)

This command specifies the locale for the operating language of the appliance.

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 Administration section 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 269](#) and [“Changing appliance names in disaster recovery configurations” on page 290](#).

Procedure

1. Start the IBM MQ Appliance web UI, and click the administration icon .
2. Select **Device > 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.

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 269](#) and [“Changing appliance names in disaster recovery configurations” on page 290](#).

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). 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.

Related reference

[“name” on page 979](#)

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 207](#). 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
```

Related reference

[“name” on page 979](#)

Configuring appliance user access

User access to the appliance is controlled by the role based management (RBM) feature.

See [“Role based management” on page 439](#) for details of how to control the authentication and authorization of appliance users.

Configuring the REST management interface

You can use the REST management interface to configure and control your appliance using REST requests.

You must configure REST management interface to enable it and define how it is accessed.

Configuring the REST management interface by using the IBM MQ Appliance web UI

You can configure the REST management interface by using the IBM MQ Appliance web UI.

About this task

You use the Administration section of the IBM MQ Appliance web UI to configure the REST management 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 the major steps that you need to complete are described.

Procedure

1. Start the IBM MQ Appliance web UI, and click the objects icon .
2. Select **Device management > REST Management interface**.
3. Ensure that **Enable administrative state** is enabled.
4. Specify the IP address (or host alias) that is used to connect to the REST management interface. Specify 0.0.0.0 to listen on all appliance interfaces.
5. Specify the port number to connect on. The port is 5554 by default for a REST interface.
6. If you want to secure your REST connection with SSL (TSL), specify a **Custom SSL server type** of **Server Profile**, and select the Custom SSL server profile (see [“Configuring certificates for the REST management interface”](#) on page 245 for instructions on creating the profile).
7. Click **Apply** to save the changes to the running configuration.

Configuring the REST management interface by using the command line

You can configure the REST management interface by using the command interface.

About this task

To configure the REST management interface from the command line, you enter rest-mgmt mode and enter the required rest-mgmt commands.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type `rest-mgmt` to enter rest-mgmt mode.
4. Type the following command to enable the REST management interface:

```
admin-state enabled
```

5. Specify the local interface that the appliance listens on for REST requests:

```
local-interface IPaddress
```

You can specify an IP address of 0.0.0.0 to listen on all appliance interfaces.

6. Specify the port number to listen on:

```
port port_number
```

The appliance listens on port 5554 by default.

7. If you want to secure REST requests and responses by using SSL (TLS), specify the SSL configuration type of `server` and supply the name of the SSL server profile that you have created for this purpose (see [“Configuring certificates for the REST management interface”](#) on page 245).

```
ssl-config-type server
ssl-server server_profile
```

8. After you configure the REST management interface, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Related concepts

[“REST Management Interface commands”](#) on page 950

Use the REST Management Interface commands to enable and modify the configuration of the REST management interface.

Configuring certificates for the REST management interface

You can configure the REST management interface to use certificates that you supply.

About this task

You use the appliance command line interface to configure the REST management interface to use your certificates.

To set up secure communication between a REST client and the REST management interface 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 you make certificate requests for an appliance, the CN part of the distinguished name must be the URL that you type to connect to the REST API. 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 REST management service listens on all of the appliance ports (`local` address set to 0.0.0.0). However, you can configure the service so that it listens on an IP address or host alias of a specific port (and so limit access to the REST management interface).

Procedure

- To upload certificates to your appliance:
 - a) Ensure that you have the following items:
 - A private key to access the appliance certificate.
 - The appliance certificate.
 - Client certificates (optional).
 - b) Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.
 - c) Log in as a user in the administrators group.
 - d) Type the following command to enter configuration mode:

```
config
```

- e) 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 489.

- To create definition objects for the appliance certificate and key:
 - a) From configuration mode, type `crypto` to enter crypto configuration mode.
 - b) 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
```

- c) Create a crypto certificate definition for the appliance:

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

For example:

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

- d) Create a crypto credential definition for the appliance:

```
idcred credential_name key_alias cert_alias
```

For example:

```
idcred RESTmguCred01 RESTmgtKey01 RESTmgmtCert01
```

- To create a crypto valcred definition for validating clients (this step is optional):
 - a) 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 RESTClientCert01 cert:///client1.cer
certificate RESTClientCert02 cert:///client2.cer
certificate RESTClientCert03 cert:///client3.cer
```

- b) 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 RESTcred01
certificate RESTClientCert01
certificate RESTClientCert02
certificate RESTClientCert03
```

- To create an SSL server profile for the appliance:
 - a) 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 RESTmgmtCred01
protocols TLSv1d2
valcred RESTcred01
request-client-auth on
require-client-auth on
send-client-auth-ca-list on
```

- To save all the changes that you have made in crypto configuration mode:
 - a) Type `exit` to leave crypto configuration mode.
 - b) Type `write mem` to save your configuration changes.
- To associate the SSL server profile with the REST management interface:
 - a) From configuration mode, type `rest-mgmt` to enter REST management interface configuration mode.
 - b) Enter the following command:

```
ssl-server SSL_Svr_Profile_name
```

For example:

```
ssl-server myappliance1
```

- To save your REST management interface configuration:
 - a) Type `exit` to leave rest-mgmt configuration mode.
 - b) Type `write mem` to save your configuration changes.
 - c) Type `exit` again to leave configuration mode.

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

[“idcred” on page 733](#)

This command creates identification credentials.

[“valcred” on page 749](#)

This command enters Validation Credentials mode.

[“key” on page 734](#)

This command creates an alias for a private key.

[“ssl-server” on page 748](#)

This command enters SSL Server Profile mode.

Configuring the appliance by using the REST management interface

You can use the REST management interface to view and alter IBM MQ Appliance configurations.

When you use the REST management interface for this purpose, you send HTTP requests to the REST interface port and receive JSON-formatted responses with a payload and indication of success or failure. You can incorporate requests into programs and so automate interaction with the appliance.

For a reference guide to the REST management interface, see [“REST management interface” on page 1000](#).

Retrieving configuration information by using REST

There are a number of major steps that are involved in retrieving configuration information from the IBM MQ Appliance by using the REST management interface.

Identify the object class

To begin retrieving the required configuration information from the appliance, first identify the specific object class that you need. The configuration root URI for as `/mgmt/config/`. To retrieve a list of all available configuration object classes on the appliance, make a request based on the following example:

```
GET https://mqhost.com:5554/mgmt/config/
```

To identify the exact formatting of the object class name, you search the received response payload. The following listing shows some fragments of the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/config/"
    }
  },
  "AccessControllist": {
    "href": "/mgmt/config/{domain}/AccessControllist"
  },
  "AuditLog": {
    "href": "/mgmt/config/{domain}/AuditLog"
  },
  "CertMonitor": {
    "href": "/mgmt/config/{domain}/CertMonitor"
  },
  "CRLFetch": {
    "href": "/mgmt/config/{domain}/CRLFetch"
  },
  "CryptoCertificate": {
    "href": "/mgmt/config/{domain}/CryptoCertificate"
  },
  "...": {
    "TimeSettings": {
      "href": "/mgmt/config/{domain}/TimeSettings"
    }
  },
  "TraceTarget": {
    "href": "/mgmt/config/{domain}/TraceTarget"
  },
  "User": {
    "href": "/mgmt/config/{domain}/User"
  },
  "UserGroup": {
    "href": "/mgmt/config/{domain}/UserGroup"
  },
  "...": {
  }
}
```

Alternatively, you can examine the URI that is displayed in the web UI browser window when an object or object list is accessed to identify the format of the object class.

Retrieve the object class list

After you identify the required object class name, you can retrieve a list of objects that exist for that class. To retrieve the list, you construct a URI of the form `/mgmt/status/domain/class_name`, replacing *domain* with the string "default" and *class_name* with the desired object class. The following request shows a URI to retrieve information from the User object class within the default domain:

```
https://mqhost.com:5554/mgmt/status/default/User
```

The User object returns the following information:

```
{
  "_links" : {
    "self" : {
      "href" : "/mgmt/config/default/User"
    }
  },
  "doc" : {
    "href" : "/mgmt/docs/config/User"
  }
},
{
  "User" : [
    {
      "name" : "admin",
      "_links" : {
        "self" : {
          "href" : "/mgmt/config/default/User/admin"
        }
      },
      "doc" : {
        "href" : "/mgmt/docs/config/User"
      }
    },
    {
      "mAdminState" : "enabled",
      "UserSummary" : "Administrator",
      "AccessLevel" : "privileged"
    }
  ],
  {
    "name" : "bob",
    "_links" : {
      "self" : {
        "href" : "/mgmt/config/default/User/bob"
      }
    },
    "doc" : {
      "href" : "/mgmt/docs/config/User"
    }
  }
},
{
  "mAdminState" : "enabled",
  "UserSummary" : "",
  "AccessLevel" : "group-defined",
  "GroupName" : {
    "value" : "Viewer",
    "href" : "/mgmt/config/default/UserGroup/Viewer"
  }
}
]
}
```

Retrieve an individual object

You can also retrieve the configuration information about a specific object, instead of retrieving the object list output in its entirety. To retrieve information about a specific object, you construct a URI of the form `/mgmt/config/domain/class_name/object_name`. You replace *domain* with the string "default", *class_name* with the required object class, and *object_name* with the name of a particular object that has been configured. For example, you could enter the following URI to retrieve just the details for the user with the ID "bob":

```
https://mqhost.com:5554/mgmt/config/default/User/bob
```

The object returns the following information:

```

{
  "_links" : {
    "self" : {
      "href" : "/mgmt/config/default/User/bob"
    },
    "doc" : {
      "href" : "/mgmt/docs/config/User"
    }
  },
  "User" : {
    "name" : "bob",
    "mAdminState" : "enabled",
    "UserSummary" : "",
    "AccessLevel" : "group-defined",
    "GroupName" : {
      "value": "Viewer",
      "href" : "/mgmt/config/default/UserGroup/Viewer"
    }
  }
}

```

Retrieve an individual object property

You can also retrieve the value of a particular property from the configuration information of a specific object. To retrieve the value of a property, you construct a URI of the form `/mgmt/config/domain/class_name/object_name/property_name`. You replace *domain* with the string "default", *class_name* with the required object class, *object_name* with the name of a particular object that has been configured, and *property_name* with the name of the property whose value you want to retrieve. For example, you could enter the following URI to retrieve the value of the `AccessLevel` property for the user with the ID "bob":

```
https://mqhost.com:5554/mgmt/config/default/User/bob/AccessLevel
```

The object returns the following information:

```

{
  "_links" : {
    "self" : {
      "href" : "/mgmt/config/default/User/bob/AccessLevel"
    },
    "doc" : {
      "href" : "/mgmt/docs/config/User/AccessLevel"
    }
  },
  "AccessLevel" : "group-defined"
}

```

Related concepts

[“REST management interface” on page 1000](#)

The REST management interface provides access to the actions and to the configuration and status resources on the appliance.

Modifying and deleting existing configurations by using REST

The steps that are involved in modifying or deleting an existing configuration on the IBM MQ Appliance depend on the level of the change that you want to make.

Modify the property-level configuration

To modify an existing property value, you overwrite the existing value with a payload that contains an updated value. To overwrite the value, first retrieve the current property value, see [Retrieve an individual object property](#).

In this example, you want to change the IP address that is assigned to Ethernet interface 4 on the appliance. First, you retrieve the current value of the `IPAddress` property of the `eth4` object that belongs to the `EthernetInterface` object class:

```
GET https://mqhost.com:5554/mgmt/config/default/EthernetInterface/eth4/IPAddress
```

You receive the following response:

```
{
  "_links" : {
    "self" : {
      "href" : "/mgmt/config/default/EthernetInterface/Eth4/IPAddress"
    },
    "doc" : {
      "href" : "/mgmt/docs/config/EthernetInterface/IPAddress"
    }
  },
  "IPAddress" : "198.51.100.1/24"
}
```

Modify the property payload that is received to remove the `_links` stanza and change the property value to the new required value. Any properties that reference other configuration objects on the appliance must also remove the embedded href link. In the case of the `IPAddress` property example, the following listing shows the modified payload:

```
{"IPAddress" : "203.0.113.10/24"}
```

After the modified payload is composed, you can overwrite the existing property value by sending an HTTP PUT request as shown in the following example. Updating the configuration on the property level allows for updating one property value per request.

```
PUT https://mqhost.com:5554/mgmt/config/default/EthernetInterface/eth4
```

After the target property is updated, a confirmation response is received:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/config/default/EthernetInterface/eth4/IPAddress"
    },
    "doc": {
      "href": "/mgmt/docs/config/EthernetInterface"
    }
  },
  "IPAddress": "property has been updated."
}
```

Modify the object-level configuration

To update multiple property values with a single request, an update on the object level is required.

To modify an existing object configuration, overwrite the existing configuration with an updated payload. To overwrite the configuration, retrieve the current configuration of the object to be modified. In the following example, the configuration for the host alias object that is named "Thur_server" is retrieved.

```
GET https://mqhost.com:5554/mgmt/config/default/HostAlias/Thur_server
```

```
{
  "_links" : {
    "self" : {
      "href" : "/mgmt/config/default/HostAlias/Thur_server"
    },
    "doc" : {
      "href" : "/mgmt/docs/config/HostAlias"
    }
  },
  "mAdminState" : "enabled",
  "UserSummary" : "The thursday server",
  "IPAddress" : "198.51.100.30"
}
```

The payload is modified to remove the `_links` stanza and amend property values as required:

```
{
  "mAdminState" : "enabled",
  "UserSummary" : "The Thurleigh server",
  "IPAddress" : "198.51.100.99"
}
```

You then put the payload that describes the `Thur_server` object back to the `HostAlias` configuration:

```
PUT https://mqhost.com:5554/mgmt/config/default/HostAlias
```

After the target object is updated, a confirmation response is received:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/config/default/HostAlias/Thur_server"
    },
    "doc": {
      "href": "/mgmt/docs/config/HostAlias"
    }
  },
  "Thur_server": "Configuration has been updated."
}
```

Delete the object-level configuration

You can delete configurations at the object level. For example, you can delete a particular host alias from the `HostAlias` configuration:

```
DELETE https://mqhost.com:5554/mgmt/config/default/HostAlias/Green_server
```

You receive confirmation that the deletion has succeeded:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/config/default/HostAlias/Green_server"
    },
    "doc": {
      "href": "/mgmt/docs/config/HostAlias"
    }
  },
  "HostAlias": {
    "value": "Configuration has been deleted."
  }
}
```

Related concepts

“REST management interface” on page 1000

The REST management interface provides access to the actions and to the configuration and status resources on the appliance.

Creating configurations by using REST

You can create a new configuration by using the REST management interface.

Compose the valid request payload

To create an object configuration, create a valid payload that contains the new configuration. To begin constructing the payload, identify the structural description of the target object, that is, the object schema. You can retrieve a schema from the metadata resource of the REST management interface by using a request of the following form:

```
GET https://mqhost.com:5554/mgmt/metadata/default/object_name
```

Where *object_name* identifies the configuration object that you want to create. For example, to retrieve the metadata resource for the host alias configuration object, you would make the following request:

```
GET https://mqhost.com:5554/mgmt/metadata/default/HostAlias
```

You receive the following response:

```
{
  "_links" : {
    "self" : {
      "href" : "/mgmt/metadata/default/HostAlias"
    },
    "doc" : {
      "href" : "/mgmt/docs/metadata/HostAlias"
    }
  },
  "object" : {
    "name" : "HostAlias",
    "uri" : "network/host-alias",
    "cli-alias" : "host-alias",
    "properties" : {
      "property" : [{
        "name" : "mAdminState",
        "type" : {
          "href" : "/mgmt/types/default/dmAdminState"
        },
        "cli-alias" : "admin-state",
        "default" : "enabled",
        ...
      },
      {
        "name" : "UserSummary",
        "type" : {
          "href" : "/mgmt/types/default/dmString"
        },
        "cli-alias" : "summary",
        "display" : "Comments",
        ...
      },
      {
        "name" : "IPAddress",
        "type" : {
          "href" : "/mgmt/types/default/dmIPHostAddress"
        },
        "required" : "true",
        "cli-alias" : "ip-address",
        "display" : "IP address",
        ...
      }
    ]
  }
}
```

You can also acquire the metadata for the object that you want to create by looking up the appliance Service-Oriented Management Interface (SOMA) schema for the configuration object. The SOMA schemas are located in the `store:///xml-mgmt.xsd` file.

From the resource metadata, you can identify the properties that are required to create the target object. You can also identify the property names to use in the payload. By using this information, you can create a proper JSON request payload. A JSON payload has the following structure:

```
{
  "object_class_name": {
    "name": "object_name",
    "property1_name": "property1_value",
    "property2_name": "property2_value",
    "property3_name": "property3_value",
    "property4_name": "property4_value"
    ...
  }
}
```

Using the information that you retrieved about the host alias configuration object, you could create the following payload:

```
{
  "HostAlias": {
```

```

    "name": "Key_server",
    "UserSummary": "Alias for Keysoe server",
    "IPAddress": "198.51.100.30",
  }
}

```

Compose the valid request URI

You can choose from two approaches to create a configuration object. Both approaches achieve the same result, but target a different URI. The first approach uses an HTTP POST request, and the second uses an HTTP PUT request. Use the POST request to create objects because a POST request results in failure if an object with the same name exists in the target domain. This approach prevents you from accidentally overwriting an existing object configuration. However, you can create an object configuration by using a PUT request. Issuing a PUT request on an existing object configuration overwrites the configuration with the values in the request payload.

The following POST request could be used to create the host alias object that is defined by the example payload:

```
POST https://mqhost.com:5554/mgmt/config/default/HostAlias
```

The following PUT request could also be used to create the host alias object that is defined by the example payload:

```
PUT https://mqhost.com:5554/mgmt/config/default/HostAlias/Key_server
```

When you send the request, both the POST and PUT requests return the same response after the object is successfully created:

```

{
  "result": "",
  "_links": {
    "self": {
      "href": "/mgmt/config/default/HostAlias"
    },
    "doc": {
      "href": "/mgmt/docs/config/HostAlias"
    }
  },
  "Key_server": "Configuration has been created."
}

```

Note that, if you repeat the POST request with the same payload, the command will fail. If you repeat the PUT request, the command will succeed.

Related concepts

[“REST management interface” on page 1000](#)

The REST management interface provides access to the actions and to the configuration and status resources on the appliance.

Configuring user access to the IBM MQ Console and the CLI

You can configure the appliance so that different users have different levels of access to the console and the CLI.

To illustrate how you can configure the appliance in this way, this topic implements the following scenario:

- Alice requires full administrative access to both appliance system settings and IBM MQ.
- Bob requires administrative access to appliance system settings but he does not require access to IBM MQ.
- Carlos requires full administrative access to the IBM MQ Console but no access to appliance system settings.
- Dave requires full administrative access to the IBM MQ Console and access to the IBM MQ CLI.

- Erin requires read-only administrative access to the IBM MQ Console so she can monitor IBM MQ and its configuration.
- Frank requires limited access to one queue manager using the IBM MQ Console.

Granting full administrative access to appliance system settings and IBM MQ

In this scenario, user Alice is granted full administrative access to the appliance and to IBM MQ.

About this task

There are two different ways that you can grant Alice the user access that she requires:

- You can create a privileged local user account for Alice
- You can add Alice as a user to a user group that grants full administrative access.

In this scenario, we use the IBM MQ Appliance web UI for all tasks, and assume that you are the admin user.

Procedure

- To create a privileged local user account, complete the following steps:

- Start the IBM MQ Appliance web UI, and click the administration icon .
- Select **Access > User Account**.
- Enter a name for the user account, in this case enter *Alice*.
- Specify an initial password for Alice.
- Select the access level **Privileged**.
- Click **Apply** to create the user account.

- To create a user group with administrative access, and add Alice to it, complete the following steps:

- Start the IBM MQ Appliance web UI, and click the administration icon .
- Select **Access > User Group**.
- Select **New**.
- Enter a name for the user group, in this case enter *Administrators*.
- Specify the following access policy in the access profile:

```
*/*/?*Access=r+w+a+d+x
```

This profile grants read, write, add, delete, and execute access to all resources on the appliance.

- Click **Apply** to create the user group
- Create a user account for Alice. Select an **Access level** of **Group defined**, and select the **Administrators** group that you just created in **User group**.
- Click **Apply** to create the user account.

Related concepts

[“User authorization, credential mapping, and access profiles” on page 423](#)

The credential mapping part of role based management authorizes appliance users to use different features on the IBM MQ Appliance.

Granting full administrative access to appliance system settings but barring access to IBM MQ

In this scenario, user Bob is granted full administrative access to the appliance, but no access to IBM MQ.

About this task

Bob's user access is configured by defining a user group and adding Bob to that group.

In this scenario, we use the IBM MQ Appliance web UI for all tasks, and assume that you are the admin user.

Procedure

To create a user group with the required access, and add Bob to it, complete the following steps:

1. Start the IBM MQ Appliance web UI, and click the administration icon .
2. Select **Access > User Group**.
3. Select **New**.
4. Enter a name for the user group, in this case enter `Appliance_admin`.
5. Specify the following four access policies in the access profile:

- Define a policy that grants full access to all appliance resources:

```
*/*/?*Access=r+w+a+d+x
```

- Define a more specific access policy that revokes authority to the IBM MQ CLI. Click **Add** and enter the following policy:

```
*/*/mq/cli?Access=NONE
```

- Define another access policy that revokes admin authority for the IBM MQ Console. Click **Add** and enter the following policy:

```
*/*/mq/webadmin?Access=NONE
```

- Define another access policy that revokes user authority for the IBM MQ Console. Click **Add** and enter the following policy:

```
*/*/mq/webuser?Access=NONE
```

You can also use the policy builder to define the access policies interactively. If you use the builder, specify the following resources:

- (all resources) (read, write, add, delete, execute privilege)
 - MQ CLI Administration (no privileges)
 - MQ Web Administration (no privileges)
 - MQ Web User (no privileges)
6. Create a user account for Bob. Select **Access > User Account** and specify the name **Bob**
 7. Select an **Access level of Group defined**, and select the **Appliance_admin** group that you just created in **User group**.
 8. Click **Apply** to create the user account.

Related concepts

[“User authorization, credential mapping, and access profiles” on page 423](#)

The credential mapping part of role based management authorizes appliance users to use different features on the IBM MQ Appliance.

Granting access to IBM MQ but barring access to the appliance system settings

In this scenario, user Carlos is granted full administrative access to the IBM MQ Console, but no access to the appliance. Dave is granted the same access as Carlos, with the addition of IBM MQ CLI access. Erin is granted read-only access to the IBM MQ Console.

About this task

Carlos, Dave, and Erin's user access is configured by defining different user groups and adding Carlos, Dave, and Erin to the required groups.

In this scenario, we use the IBM MQ Appliance web UI for all tasks, and assume that you are the admin user.

Procedure

- To create a user group with access to the IBM MQ Console, and add Carlos to it, complete the following steps:

a) Start the IBM MQ Appliance web UI, and click the administration icon .

b) Select **Access > User Group**.

c) Select **New**.

d) Enter a name for the user group, in this case enter `MQConsole`.

e) Specify the following access policies in the access profile:

- Define an access policy that enables group members to log into the IBM MQ Console. Click **Add** and enter the following policy:

```
*/*/login/web-mgmt?Access=r
```

- Define another access policy that grants read and write administrative access to all resources in the IBM MQ Console. Click **Add** and enter the following policy:

```
*/*/mq/webadmin?Access=r+w
```

- Define another access policy that grants users in the group authority to change their own password. Click **Add** and enter the following policy:

```
*/*/access/change-password?Access=x
```

You can also use the policy builder to define the access policies. If you use the builder, specify the following resources:

- Web-Mgmt (read privilege)
- MQ Web Administration (read and write privileges)
- Change User Password (execute privilege)

f) Create a user account for Carlos. Select **Access > User Account** and specify `Car1os` as the user name.

g) Select an **Access level** of **Group defined**, and in **User group** select the **MQConsole** group that you just created.

h) Click **Apply** to create the user account.

- To create a user group with access to the IBM MQ Console and the CLI, and add Dave to it, complete the following steps:

- a) Start the IBM MQ Appliance web UI, and click the administration icon .
 - b) Select **Access > User Group**.
 - c) Select **New**.
 - d) Enter a name for the user group, in this case enter MQConsoleCLI.
 - e) Specify the same policies as added to the MQConsole group previously described. In addition, add the following policies:
 - Define an access policy that enables users to log in to the appliance. Click **Add** and enter the following policy:


```
*/*/login/ssh?Access=r
```
 - Define another access policy that grants users in the group access to the IBM MQ CLI. Click **Add** and enter the following policy:


```
*/*/mq/cli?Access=x
```

You can also use the policy builder to define the access policies. If you use the builder, specify the following resources:

 - Ssh (read privilege)
 - MQ CLI Administration (execute privilege)
 - f) Create a user account for Dave. Select **Access > User Account** and specify Dave as the user name.
 - g) Select an **Access level** of **Group defined**, and in **User group** select the **MQConsoleCLI** group that you just created.
 - h) Click **Apply** to create the user account.
- To create a user group with read-only access to the IBM MQ Console and add Erin to it, complete the following steps:

- a) Start the IBM MQ Appliance web UI, and click the administration icon .
- b) Select **Access > User Group**.
- c) Select **New**.
- d) Enter a name for the user group, in this case enter MQConsoleReadOnly.
- e) Specify the following policies:
 - Define an access policy that enables group members to log into the IBM MQ Console. Click **Add** and enter the following policy:


```
*/*/login/web-mgmt?Access=r
```
 - Define another access policy that grants read administrative access to all resources in the IBM MQ Console. Click **Add** and enter the following policy:


```
*/*/mq/webadmin?Access=r
```
 - Define another access policy that grants users in the group authority to change their own password. Click **Add** and enter the following policy:


```
*/*/access/change-password?Access=x
```

You can also use the policy builder to define the access policies. If you use the builder, specify the following resources:

- Web-Mgmt (read privilege)
- MQ Web Administration (read privilege)

- Change User Password (execute privilege)
- f) Create a user account for Erin. Select **Access > User Account** and specify Erin as the user name.
- g) Select an **Access level** of **Group defined**, and in **User group** select the **MQConsoleReadonly** group that you just created.
- h) Click **Apply** to create the user account.

Related concepts

“User authorization, credential mapping, and access profiles” on page 423

The credential mapping part of role based management authorizes appliance users to use different features on the IBM MQ Appliance.

Granting limited access to a queue manager

In this scenario user Frank is granted access to use the IBM MQ Console to view details about a single queue manager.

About this task

There are two stages to configuring a user to have access to a single queue manager, and no other parts of the IBM MQ or appliance configurations.

Firstly you create a user group that gives user access to the IBM MQ Console and add Frank to that group. You use the IBM MQ Appliance web UI to complete this stage.

Then you create a messaging user of the same name (Frank) so that MQ authorities can be granted to Frank by using the MQ object authority manager (OAM). You use the IBM MQ command line, MQCLI, to complete this stage.

Procedure

- To create a user group with access to the IBM MQ Console, and add Frank to it, complete the following steps:

- a) Start the IBM MQ Appliance web UI, and click the administration icon .
- b) Select **Access > User Group**.
- c) Select **New**.
- d) Enter a name for the user group, in this case enter MQConsoleLimited.
- e) Specify the following access policies in the access profile:

- Define an access policy that enables group members to log into the IBM MQ Console. Click **Add** and enter the following policy:

```
*/*/login/web-mgmt?Access=r
```

- Define another access policy that grants group members the required permission to access IBM MQ in the IBM MQ Console. Click **Add** and enter the following policy:

```
*/*/mq/webuser?Access=x
```

- Define another access policy that grants users in the group authority to change their own password. Click **Add** and enter the following policy:

```
*/*/access/change-password?Access=x
```

You can also use the policy builder to define the access policies. If you use the builder, specify the following resources:

- Web-Mgmt (read privilege)
- MQ Web User (execute privilege)

- Change User Password (execute privilege)
- f) Create a user account for Frank. Select **Access > User Account** and specify Frank as the user name.
- g) Select an **Access level** of **Group defined**, and in **User group** select the **MQConsoleLimited** group that you just created.
- h) Click **Apply** to create the user account.
- To define a messaging user, complete the following steps:
 - a) Log into the appliance command line, and enter the MQ CLI:

```
mqa# mqcli
mqa (mqcli)#
```

- b) Create the messaging user Frank:

```
mqa (mqcli)# usercreate -u Frank
```

You do not need to specify a password because the appliance user password is used to log in to the IBM MQ Console. See [“Administering messaging users”](#) on page 347 for more information about messaging users.

- c) You must now run MQ authority commands to give Frank the required access. You can define the access by using MQSC, and you can grant access directly to Frank (you could also define a messaging group, add Frank to it, and grant access to that group). Assuming Frank only wants to display information about the queue manager QM1 and the queues defined on it, run the following MQSC commands to grant Frank access to the IBM MQ Console to display QM1 and associated queues:

```
mqa (mqcli)# runmqsc QM1
5724-H72 (C) Copyright IBM Corp. 1994, 2020.
Starting MQSC for queue manager QM1.
SET AUTHREC PROFILE(SYSTEM.ADMIN.COMMAND.QUEUE) OBJTYPE(Queue) PRINCIPAL('Frank')
AUTHADD(put)
SET AUTHREC PROFILE(SYSTEM.REST.REPLY.QUEUE) OBJTYPE(Queue) PRINCIPAL('Frank')
AUTHADD(put,GET,INQ,BROWSE)
SET AUTHREC OBJTYPE(QMGR) PRINCIPAL('Frank') AUTHADD(DSP)
SET AUTHREC PROFILE(**) OBJTYPE(Queue) PRINCIPAL('Frank') AUTHADD(DSP)
```

You could use the IBM MQ Console instead of runmqsc to define the MQ authorities for Frank, if required.

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 311.

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 644.

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 417.

Configuring uniform clusters - appliance-specific considerations

You can configure uniform clusters on the appliance, but there are some considerations specific to the appliance.

For a full description of IBM MQ uniform clusters and how to set them up, see [Uniform clusters](#) in the IBM MQ documentation.

Both the initial creation of a uniform cluster and subsequently keeping the configuration between the uniform cluster members identical can be simplified by using the automatic configuration and automatic clustering support. When using this capability, one configuration file describes the cluster and another represents the MQSC configuration to apply to all queue managers in the uniform cluster. On each queue manager restart, the configuration is reapplied and the cluster automatically forms. Full instructions for creating these files are provided in [Creating a uniform cluster](#) in the IBM MQ documentation.

On the appliance both the auto-cluster file (with the `.ini` suffix) and the auto-config file (with the `.mqsc` suffix) must be located in the `mqbackup:` directory on the appliance.

You specify these files when you use the `crtmqm` command to create a uniform cluster queue manager.

If you are deploying a uniform cluster on a high availability (HA) or disaster recovery (DR) configuration, then you must take steps to ensure that all the nodes have the same auto-config and auto-cluster files in their `mqbackup:` directory.

When you back up uniform cluster queue managers from your appliance, you must take explicit steps to back up the auto-config and auto-cluster files from the `mqbackup:` directory.

For information about moving files to and from the appliance, see [“Managing files by using the IBM MQ Appliance web UI”](#) on page 402 and [“Managing files by using the REST management interface”](#) on page 405.

Related reference

[“crtmqm \(create queue manager\)”](#) on page 578
Create a queue manager.

[“runmqsc \(run MQSC commands\)”](#) on page 606
Run MQSC commands on a queue manager.

[“setmqini \(set mq ini\)”](#) on page 620

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. Set or delete a diagnostic message definition in the `mqsc.ini` file.

Configuring environment variables

You can configure environment variables either globally or for a specific queue manager.

For a list of IBM MQ environment variables, see [Environment variables](#) in the IBM MQ documentation.

Some environment variables are not applicable, or have limited usefulness on the IBM MQ Appliance. In particular, the following environment variables are never permitted (and are rejected by the `setmqvar` command)

- JAVA
- MQ_JAVA_DATA_PATH
- MQ_INSTALLATION_PATH
- PATH
- MQCHLLIB
- MQCHLTAB
- MQS_KEYSTORE_CONF

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
```

Related reference

[“setmqvar \(set mq variable\)” on page 623](#)

Add or remove an environment variable for the appliance or for a specified queue manager.

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
```

Related reference

[“setmqvar \(set mq variable\)” on page 623](#)

Add or remove an environment variable for the appliance or for a specified queue manager.

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
```

Related reference

[“setmqvar \(set mq variable\)” on page 623](#)

Add or remove an environment variable for the appliance or for a specified queue manager.

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:

```
dspmqlvar
```

- To view a specific global environment variable, enter the following command:

```
dspmqlvar -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:

```
dspmqlvar -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:

```
dspmqlvar -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:

```
dspmqlvar -k MQSSLRESET
```

Related reference

[“dspmqlvar \(display mq variables\)” on page 619](#)

Display environment variables set for a specified queue manager.

Configuring IBM MQ Advanced Message Security

IBM MQ Advanced Message Security (IBM MQ AMS) is a component of IBM MQ that provides a high level of protection for sensitive data flowing through the IBM MQ network, while not impacting the end applications.

A full description of IBM MQ Advanced Message Security are given in the IBM MQ documentation, see [IBM MQ Advanced Message Security](#).

To implement IBM MQ AMS on the appliance you must use the **runmqsc** commands to manipulate security policies for individual queue managers. Specifically, you use the following commands:

- SET POLICY

- DELETE POLICY
- DISPLAY POLICY

Note: You cannot use the IBM MQ control commands `setmqsp1` or `dspmqsp1` to work with security policies on the appliance.

Related reference

[“runmqsc \(run MQSC commands\)” on page 606](#)
Run MQSC commands on a queue manager.

Related information

[SET POLICY](#)

[DELETE POLICY](#)

[DISPLAY POLICY](#)

Configuring MCA interception

If you are using AMS to provide message-level security on a queue manager, you might want to configure MCA interception on particular server-connection channels.

MCA interception is used to implement a different message security policy for particular clients. Without MCA interception, certificates must be distributed to clients to enable them to encrypt or decrypt messages. With MCA interception configured for a channel, encryption and decryption are performed by the queue manager. Required certificates and keys must be held in the queue manager key repository (which is created automatically with the queue manager and managed with the certificate commands, see [“Queue manager security management commands” on page 629](#)). Messages in flight over the channel have no message-level encryption (but are usually protected by channel-level security, for example, TLS).

MCA interception is implemented for one of the following reasons:

- To operate in a situation where it is undesirable or not possible to distribute certificates to IBM MQ clients.
- To ensure messages are encrypted while stored on the appliance (although note that such data is still vulnerable if a disk is stolen from the appliance, or there is a malicious administrator).

You must not use MCA interception on channels that you use message-level encryption on, because this causes double-encryption.

On the appliance, you specify MCA interception for a server-connection channel using the following command:

```
setamschl -m QMgrName -n Channel_Name -c Certificate_Label
```

You can view the MCA intercept configuration of a queue manager, or specific server-connection channel, by using the following command:

```
dspamschl -m QMgrName [-n Channel_Name]
```

For more information about AMS and MCA interception, see [Message Channel Agent \(MCA\) interception in the IBM MQ documentation](#).

Related reference

[“setamschl \(set AMS channel\)” on page 643](#)

Set up MCA interception on a specific server-connection channel on a queue manager previously configured with AMS.

[“dspamschl \(set AMS channel\)” on page 637](#)

Display information about MCA interception on queue managers.

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 16](#).

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” on page 267](#).
2. Create an HA group for the appliances. For more information, see [“Configuring the high availability group” on page 271](#).
3. Create HA queue managers on the appliances. For more information, see [“Configuring high availability queue managers” on page 275](#).

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 can define a single, floating IP address that is used to connect to a queue manager on either of the appliances in the HA group (see [“Specifying a floating IP address for a queue manager” on page 277](#)). Alternatively, you can manually configure your client applications to attempt to connect to the HA appliances using their static IP addresses in your preferred order. You configure this by using the usual MQ channel connection mechanisms. For guidance on how to connect clients to queue managers, 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](#). 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.

Note:  Advanced users can, if required, configure HA using different interfaces to the default eth21, eth17, and eth13. For more information, see [“Configuring custom HA replication interfaces” on page 269](#).

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:

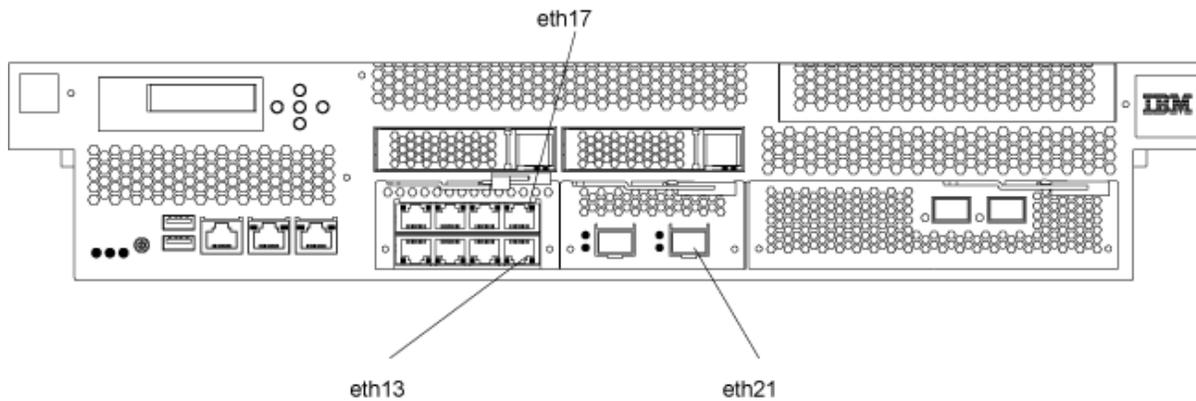


Figure 57. Ethernet ports on M2000 appliance

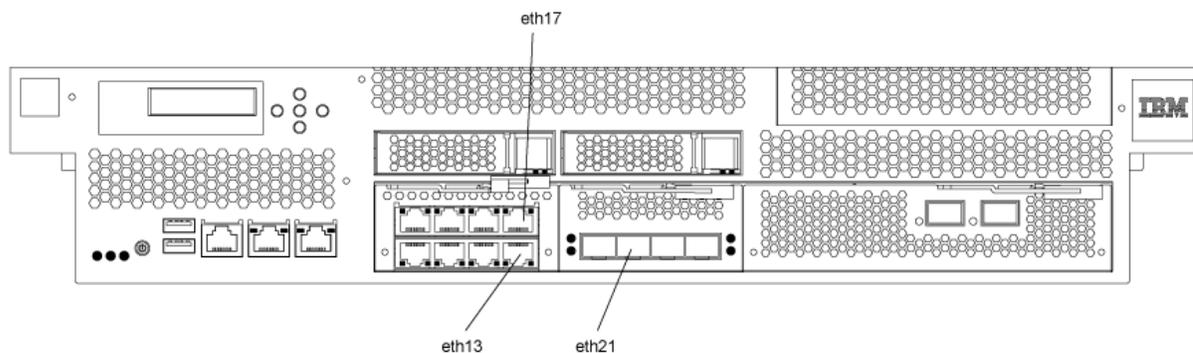


Figure 58. Ethernet ports on M2001 appliance

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 217](#). 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 242. 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 272](#).

Related concepts

[“Network requirements for high availability” on page 35](#)

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

[“Example network set up for HA configuration” on page 281](#)

The example shows the network configuration for an high availability implementation.

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 280](#), and then removing the HA group itself, see [“Deleting a high availability group” on page 274](#).
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 218](#).
3. Re-create the HA configuration by configuring the queue managers as described in [“Creating a high availability group” on page 272](#) and [“Creating a high availability queue manager” on page 275](#).

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 280](#), and then removing the HA group itself, see [“Deleting a high availability group” on page 274](#).
2. Allocate new appliance names as required, see [“Configuring the appliance name by using the command line” on page 242](#).
3. Re-create the HA configuration by configuring the queue managers as described in [“Creating a high availability group” on page 272](#) and [“Creating a high availability queue manager” on page 275](#).

v 9.1.0 Configuring custom HA replication interfaces

Advanced users can, if required, configure different replication interfaces to the default for high availability (HA) configurations.

About this task

By default, the eth21 interface is used for data replication in an HA configuration, but you can specify a different interface, if required. You can configure different interfaces on each appliance, for example, you could specify eth23 as replication interface on one appliance, and eth22 on the other. You should choose an interface of at least 10 Gb; the performance of a 1 Gb replication link is likely to be poor.

If you have added disaster recovery to your HA configuration, you can also specify a different interface for DR replication to the default eth20 interface. If you do have DR in your HA configuration, both HA appliances must use the same interface for HA replication.

Procedure

- To configure a non-default HA replication interface:
 - a) Cable your custom interface. For example, if you intended to use eth23 for HA replication on both appliances, you should connect eth23 on both appliances in your HA group, either by a direct cable or a switch.
 - b) Enter the IBM MQ administration mode on both appliances by entering the following command:
`mqcli`
 - c) Set up your custom interface by using the following command on both appliances:

```
sethalink -r IP_address
```

Where *IP_address* is the IP address of the interface to be used for replication. You must specify the IP address using ip v4 dotted decimal notation (for example, "192.0.2.8"). The two IP addresses (one for each appliance) should be in the same subnet. The interface that you specify on each appliance must meet the following requirements:

- The IP address specified must map to a valid local interface or aggregation interface (you cannot specify a VLAN interface).
- The local interface must be at least a 10 Gb link, otherwise you are warned.

The command fails if the HA configuration is already active.

- To check the configuration of your HA replication link, enter the following command:

```
dsphalink
```

The current settings for HA replication interface are displayed, for example:

```
Replication:      192.0.2.8 (1o)
The dsphalink command succeeded.
```

- To revert to the default configuration of using eth21 for the HA replication link:
 - a) Cable your eth21 interfaces on each appliance together, either directly or using a switch.
 - b) Enter the following command:

```
sethalink -d
```

The command fails if the HA configuration is already active.

- To configure a non-default DR replication link for your HA configuration:
 - a) Make the necessary hardware configuration to connect the appliances in the HA group to the DR appliance.
 - b) Enter the IBM MQ administration mode on both HA appliances by entering the following command:
`mqcli`
 - c) Set up the custom interface by entering the following command on both HA appliances:

```
setdrlink -r IP_address
```

Where *IP_address* is the IP address of the interface to be used for replication. You must specify the IP address using ip v4 dotted decimal notation (for example, "192.0.2.8"). The interface that you specify on each appliance must meet the following requirements:

- The IP address specified must map to a valid local interface or aggregation interface.
- The local interface must be at least a 10 Gb link, otherwise you are warned.
- You must select the same interface for both appliances in your HA group.

The command sets a new interface from that moment forward, so that any **crtdrprimary** and **crtdrsecondary** commands set up a DR configuration using the newly specified interface. Any

queue managers configured to use DR before this point will continue to use the default interface, rather than the new one, unless you remove and recreate the DR configuration for those queue managers.

- To check the configuration of your DR replication link, enter one of the following commands:
 - To view the DR replication link for a particular queue manager:

```
dspdlink -m QueueManager
```

The IP address and interface name currently configured for that queue manager are displayed, for example:

```
QM1          10.120.5.42 (eth21)
The dspdlink command succeeded.
```

- To view the default setting, and the settings for existing queue managers:

```
dspdlink
```

Fore example:

```
Default      10.120.10.100 (eth20)
QMGR         DR Interface (Name)
----         -
QM1          10.120.5.42 (eth21)
QM2          10.120.6.42 (eth22)
The dspdlink command succeeded.
```

- To revert to the default configuration for the DR replication link:

```
setdlink -r
```

Related reference

[“sethlink” on page 678](#)

Specify a non-default link to be used for HA replication.

[“setdlink” on page 685](#)

Specify a non-default link to be used for DR replication in an HA configuration.

[“dsphalink” on page 672](#)

Display details of the link configured for HA replication.

[“dspdlink” on page 683](#)

Display details of the link configured for DR replication in an HA group.

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 267](#)

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 app11 and app12 where a new, unique key is generated for communication between the appliances. The HA group primary interface of

app12 has the IP address 192.0.2.8; the HA group primary interface of app11 has the IP address 192.0.2.7.

The following command is run from app11:

```
prepareha -s AGEW1823510HH -a 192.0.2.8
```

The following command is run from app12:

```
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 275](#).

Related reference

[“crthagrp \(create HA group\)” on page 669](#)

Creates a high availability (HA) group of appliances.

[“dlthagrp \(delete HA group\)” on page 672](#)

Deletes a high availability (HA) group.

[“prepareha \(prepare appliance for HA group\)” on page 674](#)

Prepares an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.

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 **dsphagr** command on the command line.

About this task

The **dsphagr** 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:

```
dsphagr
```

3. Optional: Exit the IBM MQ administration mode by entering the following command:

```
exit
```

Related reference

[“dsphagr \(display HA group status\)” on page 670](#)

Displays the status of the appliances in the high availability (HA) group.

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 view the **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
```

Related reference

[“dlthagrp \(delete HA group\)” on page 672](#)

Deletes a high availability (HA) group.

[“crthagrp \(create HA group\)” on page 669](#)

Creates a high availability (HA) group of appliances.

[“prepareha \(prepare appliance for HA group\)” on page 674](#)

Prepares an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.

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 view the **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.

The following topics give instructions on how to configure high availability queue managers using the command line interface. You can also use the IBM MQ Console, see [“Working with queue managers” on page 311](#).

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 272](#).

V 9.1.4 The auto-config and auto-cluster files for a queue manager that is part of a uniform cluster must be copied to the mqbackup: directory for appliances in the high availability group. Any changes that you subsequently make to these files must be replicated manually to appliances in the group (see [“Configuring uniform clusters - appliance-specific considerations” on page 261](#)).

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 do not need to use the **strmqm** command to start it.
 - You can use other **crtmqm** parameters in the command. For more information about the available parameters, see [“crtmqm \(create queue manager\)” on page 578](#).
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
```

Related reference

[“crtmqm \(create queue manager\)” on page 578](#)

Create a queue manager.

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
```

Related reference

[“sethagrp \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

Specifying a floating IP address for a queue manager

You can optionally specify a floating IP address for a high availability (HA) queue manager such that an application can connect whichever appliance the queue manager is running on.

Before you begin

Both appliances in the HA pair must be active when you specify a floating IP address for an HA queue manager.

About this task

If you specify a floating IP address for a queue manager, an application can use that address to connect to a queue manager regardless of which appliance in the HA pair the queue manager is actually running on.

You can define only one floating IP address for IBM MQ traffic on a queue manager, so you can only run the **sethaint** command once for each queue manager.

When you specify the floating IP address for IBM MQ traffic, you also specify the local interface that it can be reached on (for example, eth22). This interface must be a physical interface that exists on both appliances, and each interface must have a static IP address configured. You can also specify an aggregate interface that exists with the same name on each appliance. The name of the aggregate interface must obey certain rules to be used for this purpose, see [“link-aggregation” on page 803](#).

The floating IP address must be a valid IPv4 address that is not already defined on either appliance, and it must belong to the same subnet as the static IP addresses defined for the local interface.

Note: If you make any changes to the configuration of the physical interface that is connected to the floating IP, then the HA queue manager that uses that floating IP will stop and cannot be restarted. This also applies if you change any aggregated interfaces. You must undo the changes to the interface before you can restart the HA queue manager by using the **strmqm** command. Alternatively, you can delete the queue manager when in this situation.

Procedure

1. Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```

2. Enter the following command to add the floating IP address to the queue manager:

```
sethaint -m queue_manager_name -a -f floating_IP -l local_interface
```

Where:

queue_manager_name

Is the queue manager that the floating IP address applies to.

floating_IP

Is the floating IP address in IPv4 format.

local_interface

Specifies the name of the local interface that is used to connect to the queue manager on the two appliances in the HA group. For example, eth22. You can also specify an aggregate interface name, provided that the name obeys certain naming rules (see [“link-aggregation” on page 803](#)).

3. Optional: Exit the IBM MQ administration mode by entering the following command:

```
exit
```

Example

The following example shows the floating IP address 192.0.2.15 being allocated for queue manager QM1 and associated with the local interface eth22:

```
sethaint -m QM1 -a -f 192.0.2.15 - l eth22
```

What to do next

After you have defined a floating IP address for a queue manager, you can bind that address to a listener or a channel.

For example, the following MQSC command binds a listener named "listy" to the floating IP address 192.0.2.15:

```
DEFINE LISTENER(listy) IPADDR(192.0.2.15)
```

The following MQSC command binds a sender channel named "sendy" to the floating IP address 192.0.2.15:

```
DEFINE CHANNEL(sendy) CHLTYPE(SDR) LOCLADDR(192.0.2.15)
```

If a queue manager with a floating IP address has to create a channel automatically, and that channel has to use the floating IP address for outbound communication, you must set the MQ_LCLADDR environment variable to the floating IP address. For example, the queue manager might create a CLUSSDR channel from a CLUSRCVR channel definition received from another queue manager. The following MQCLI command sets the MQ_LCLADDR to the floating IP address 192.0.2.15:

```
setmqvar -m QMgrName -k MQ_LCLADDR -v 192.0.2.15
```

You can use the local-address parameter of the **test tcp-connection** command to test your floating IP address, see [“test tcp-connection” on page 717](#).

Related reference

[“sethaint” on page 677](#)

Specify a floating IP address for a high availability (HA) queue manager, or delete an existing floating IP address.

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.

Remote appliance(s) unavailable

This status means that the replication connection to the other appliance has been lost.

Partitioned

Queue manager data on the appliances is out of step, and cannot be automatically resolved.

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.

Inconsistent

The status is displayed on a secondary appliance during the initial synchronization of a queue manager if connection has been lost and synchronization was interrupted. The secondary appliance cannot provide high availability functionality until the initial synchronization has completed.

Resource failed

The physical interface associated with the floating IP used by the queue manager has been reconfigured. The queue manager is stopped and cannot be restarted until the original configuration is restored. This can also apply if an aggregated interface associated with the floating IP is changed.

This situation can also arise if the queue manager has been configured for disaster recovery and the eth20 interface has been reconfigured.

To recover, you must restore the original interface configuration, or remove the disaster recovery configuration. You can also delete the queue manager.

V 9.1.0 Migration Pending

When you have upgraded the appliance from Version 8.0 or Version 9.0 to Version 9.1, high availability queue managers display this status until you run the `sethagrp -u` command to upgrade them, and the other appliance is upgraded to Version 9.1 too. .

- 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 last synchronized. This information is shown when the status is `Partitioned` or `Inconsistent`.
- The time and date that the data was last in sync. This is displayed whenever data is out of sync.

Procedure

- To view the HA status of a queue manager by using the command line interface:
 - a) Enter the IBM MQ administration mode by entering the following command:


```
mqcli
```
 - b) 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.
 - c) 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:
 - a) Open the console and find the widget that displays the queue manager.
 - b) Select the queue manager in the widget and select the properties icon from the toolbar  .
 - c) In the properties window, click on the **High availability status** section to open it.

Related reference

[“status” on page 881](#)

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.

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 stand-alone queue manager by using the **sethagr** 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 294](#).

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 stand-alone queue manager:

```
sethagr -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:

```
sethagr -e HAQM1
```

Related reference

[“sethagr \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“strmqm \(start queue manager\)” on page 610](#)

Start a queue manager.

[“dspmq \(display queue managers\)” on page 585](#)

Display information about queue managers.

[“status” on page 881](#)

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.

Configuring error reporting for high availability

Certain events in a high availability configuration cause messages to be written to the IBM MQ system log. You can write these messages to an appliance log target.

High availability messages are written to the system log, rather than individual queue manager logs, because notifications are made even when the queue manager is not running. You can interrogate the log to retrieve high availability messages, and set up log targets to receive messages (see [“IBM MQ diagnostic message services” on page 505](#) and [“Appliance log targets” on page 516](#)).

For a guide to interpreting high availability messages, see [“HA notification examples” on page 540](#).

Note: You can interrogate the status of an HA group of appliances, or an HA queue manager at any time by using the command line or the console, see [“Viewing the status of appliances in a high availability group” on page 273](#) and [“Viewing the status of a high availability queue manager” on page 278](#).

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).

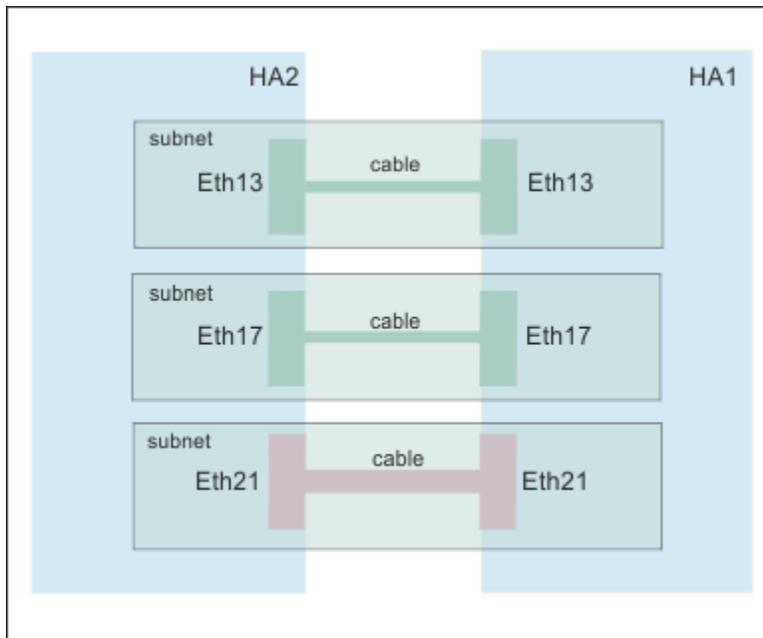


Figure 59. Example HA/DR network configuration

The following tables show how the Ethernet ports are configured on the two HA appliances, HA1 and HA2.

<i>Table 23. Ethernet ports on appliance HA1</i>			
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

<i>Table 24. Ethernet ports on appliance HA2</i>			
Ethernet port	Example IP address	Port	Description
eth13	192.0.10.12	5404, 5405 for heartbeat, 2222 for configuration	HA primary group interface
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

Related concepts

[“Network requirements for high availability” on page 35](#)

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

Related tasks

[“Configuring the hardware for high availability” on page 267](#)

You must physically connect the appliances together before you can configure the IBM MQ Appliance for high availability (HA).

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 18](#) 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. Do not stop the HA queue manager, the ending of the queue manager is ended by the underlying high availability system. (If you manually end the HA queue manager, the **crtdrprimary** command might hang, in which case contact IBM Support.)
3. Specify that the queue manager is the primary instance in a disaster recovery configuration and optionally include a floating IP address that can be used by either of the appliances in the HA pair:

```
crtdrprimary -m queue_manager_name -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.

V 9.1.5 -f floatingIP

If your HA appliances have static IP addresses assigned to the replication interface (eth20) that are in the same subnet, you can specify a floating IP address. 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. 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.

Note: For versions earlier than IBM MQ Appliance Version 9.1.5 the specification of the floating IP address is mandatory, and appliances in the HA group must be in the same subnet.

In Version 9.1.5, you can currently configure only one HA/DR queue manager with no floating IP address, contact IBM Support for help with overcoming this limitation.

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
```

If you specified a floating IP, then the **-i** parameter is the floating IP, if you omitted the **-f floating_IP** argument, then the **-i** parameter is the static IP address of the replication interface of the appliance that you ran the **crtdrprimary** command on (or the appliance that is the preferred location for the HA queue manager, if that is different). If the queue manager fails over to the other appliance in the HA pair, that appliance identifies itself to the DR appliance so that replication can continue.

- Copy the **crtldrsecondary** 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.

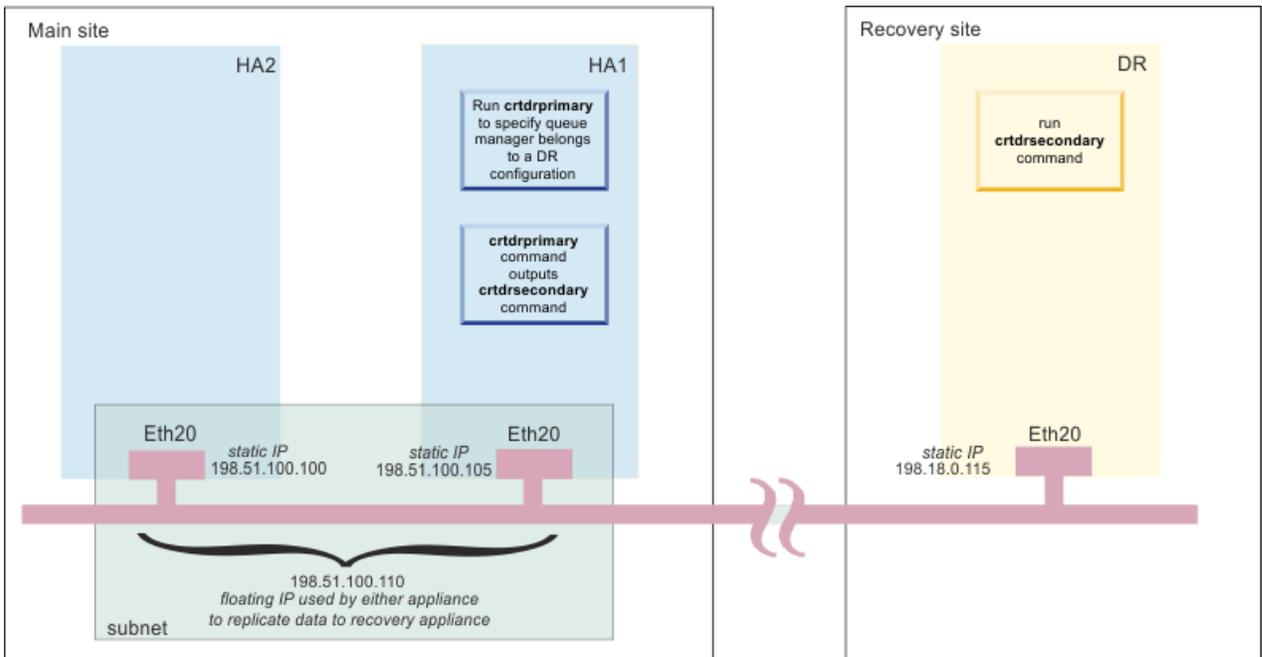


Figure 60. Configuring an HA queue manager for disaster recovery (using example IP addresses in the same subnet)

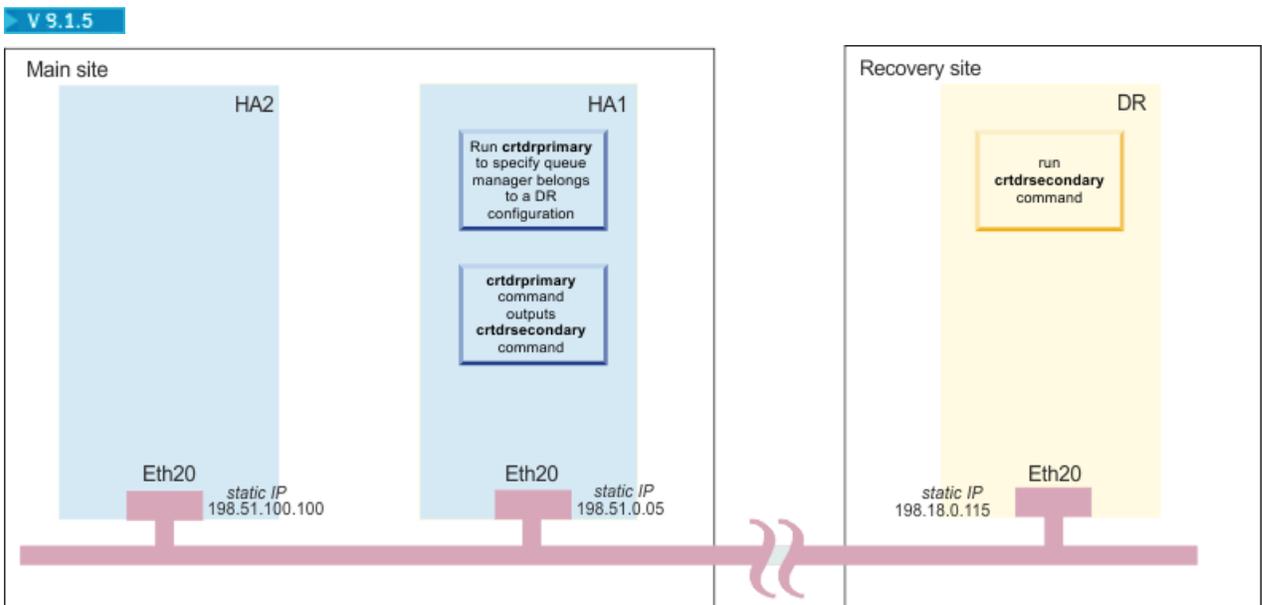


Figure 61. Configuring an HA queue manager for disaster recovery

Related concepts

“Planning network connections” on page 38

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

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.

V 9.1.5 In this example a floating IP is used to replicate data between the HA appliance currently running the queue manager and the DR appliance. The use of a floating IP is optional, and is restricted to the situation where both HA appliances replication (eth20) interfaces are in the same subnet. (For versions earlier than IBM MQ Appliance Version 9.1.5 the specification of the floating IP address is mandatory.)

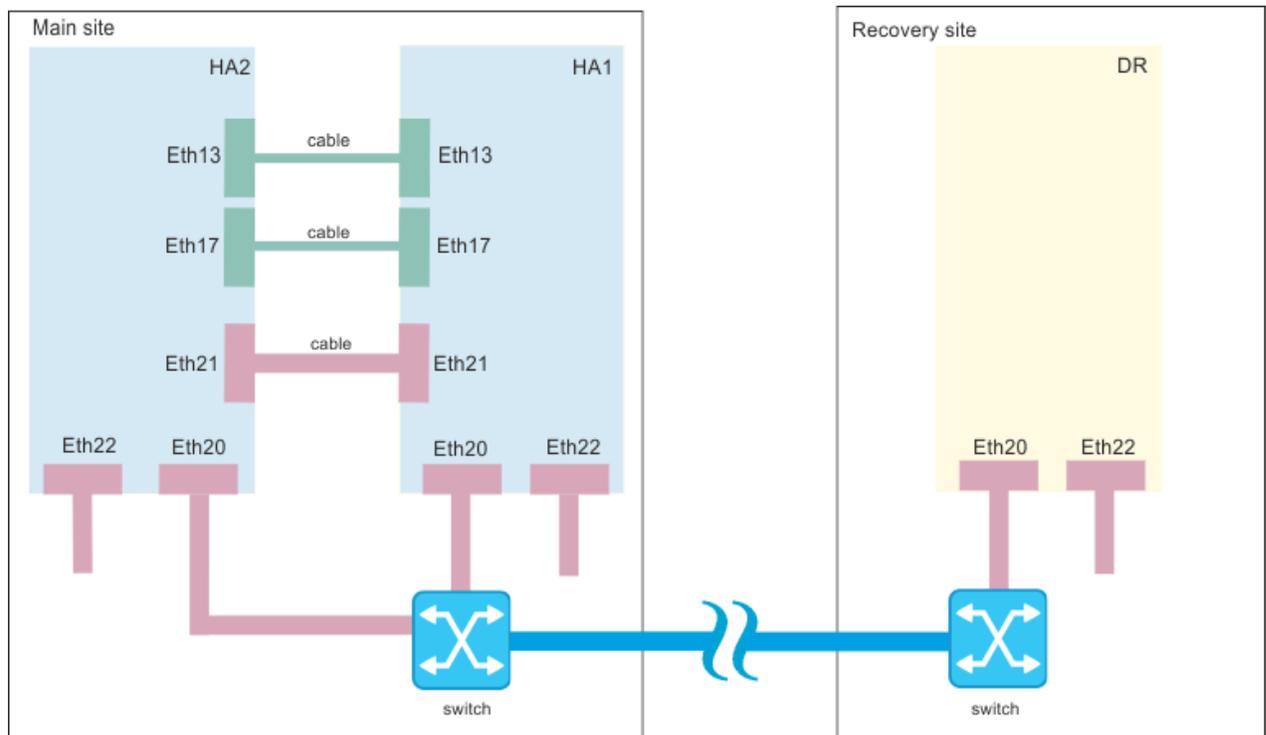


Figure 62. Example HA/DR network configuration

The following tables show how the Ethernet ports are configured on the two HA appliances, HA1 and HA2, and the DR appliance, DR.

Ethernet port	IP address	Description
eth13	192.0.10.11	HA primary group interface
eth17	192.0.20.11	HA group alternative interface
eth20	198.51.100.70	Static IP address configured for DR port
-	198.51.100.19	Floating IP address used by appliance HA1 or HA2 for replication with DR appliance
eth21	192.0.30.11	HA replication interface

<i>Table 25. Ethernet ports on appliance HA1 (continued)</i>		
Ethernet port	IP address	Description
eth22	203.0.113.0	Data interface

<i>Table 26. Ethernet ports on appliance HA2</i>		
Ethernet port	IP address	Description
eth13	192.0.10.12	HA primary group interface
eth17	192.0.20.12	HA group alternative interface
eth20	198.51.100.12	Static IP address configured for DR port
-	198.51.100.19	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

<i>Table 27. Ethernet ports on appliance DR</i>		
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 stand-alone 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:

```
dltdprimary -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.

```
sethagr -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 272](#).
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.

```
sethagr -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:

```
sethappreferred 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 282](#).

Related reference

[“endmqm \(end queue manager\)” on page 597](#)

Stop a queue manager.

[“dlthagrp \(delete HA group\)” on page 672](#)

Deletes a high availability (HA) group.

[“sethagr \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“strmqm \(start queue manager\)” on page 610](#)

Start a queue manager.

[“prepareha \(prepare appliance for HA group\)” on page 674](#)

Prepares an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.

[“crthagrp \(create HA group\)” on page 669](#)

Creates a high availability (HA) group of appliances.

[“sethappreferred \(set HA preferred appliance\)” on page 679](#)

Sets a preferred appliance in the high availability (HA) group for a queue manager to run on.

[“clearhappreferred \(clear HA preferred appliance\)” on page 679](#)

Clears the preferred appliance for a queue manager in a high availability (HA) group.

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 19. For more information about disaster recovery for a high availability pair, see [“Disaster recovery for a high availability configuration”](#) on page 18.

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:

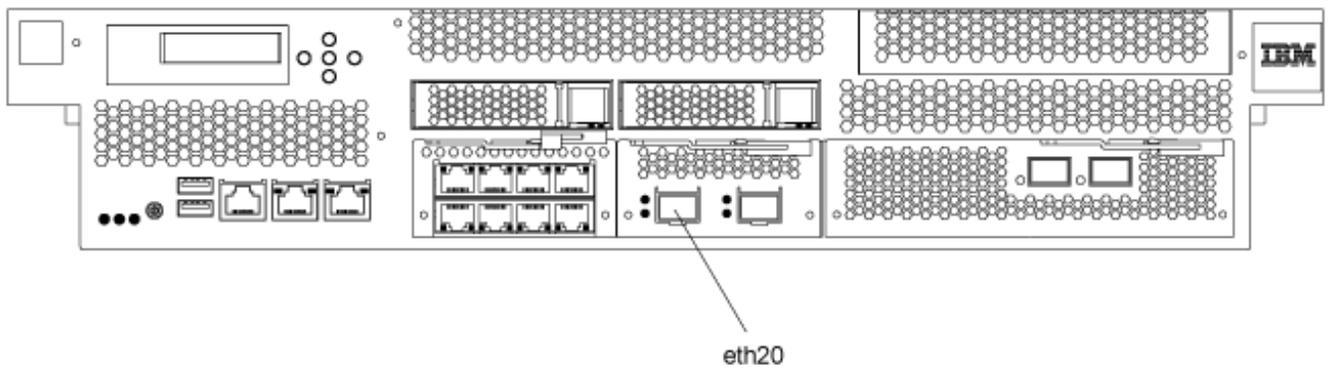


Figure 63. Eth20 on an M2000 appliance

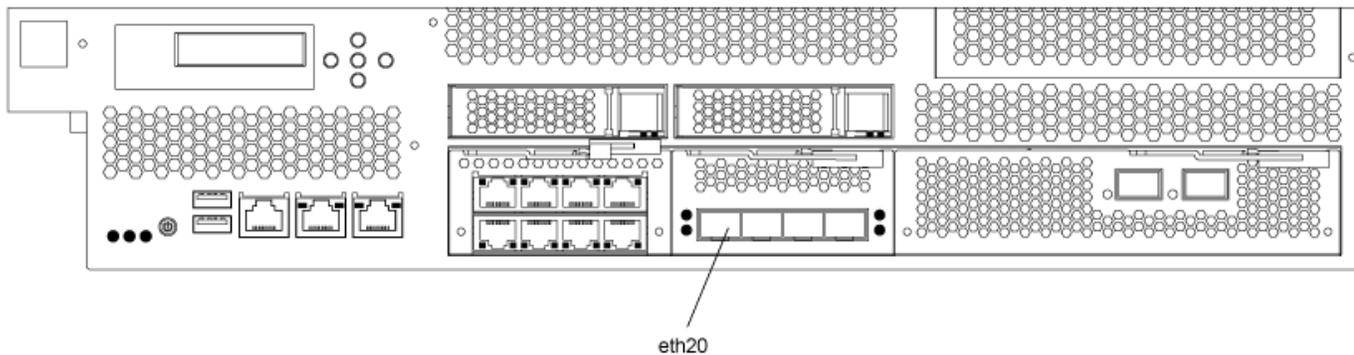


Figure 64. Eth20 on an M2001 appliance

Note: V9.1.0 Advanced users can, if required, configure DR using a different replication interface to the default eth20. For more information, see [“Configuring custom DR replication interfaces”](#) on page 290.

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 217. 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 774).
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 242. 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. Other changes to the eth20 interface on either appliance could stop high availability (HA) queue managers that also have DR enabled from running.

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 294.
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 218.
3. Re-create the DR configuration by configuring the queue managers as described in [“Configuring queue managers for disaster recovery by using the command line”](#) on page 291.

You can also use the IBM MQ Appliance web UI to perform these operations. Follow the related links for details.

If you find that your HA/DR queue managers do not run, and you receive an error message indicating that the queue manager has failed due to a problem with a resource, then you must undo the changes that were made to eth20 before you can restart the queue managers by using the **strmqm** command. Alternatively, you can delete the queue manager or remove the disaster recovery configuration from it.

Related tasks

[“Removing a queue manager from a disaster recovery configuration by using the IBM MQ Console”](#) on page 294

You can remove a queue manager from a disaster recovery configuration by using the IBM MQ Console.

[“Configuring Ethernet interfaces by using the IBM MQ Appliance web UI”](#) on page 217

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

[“Configuring queue managers for disaster recovery by using the IBM MQ Console”](#) on page 293

You set up a disaster recovery configuration by setting up a primary and secondary queue manager.

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 294.
2. Allocate new appliance names as required, [“Configuring the appliance name by using the command line”](#) on page 242.
3. Re-create the DR configuration by configuring the queue managers as described in [“Configuring queue managers for disaster recovery by using the command line”](#) on page 291.

You can also use the IBM MQ Appliance web UI to perform these operations. Follow the related links for details.

Related tasks

[“Removing a queue manager from a disaster recovery configuration by using the IBM MQ Console”](#) on page 294

You can remove a queue manager from a disaster recovery configuration by using the IBM MQ Console.

[“Configuring the appliance name by using the IBM MQ Appliance web UI”](#) on page 242

You can configure the appliance name by using the IBM MQ Appliance web UI.

[“Configuring queue managers for disaster recovery by using the IBM MQ Console”](#) on page 293

You set up a disaster recovery configuration by setting up a primary and secondary queue manager.

V 9.1.0 Configuring custom DR replication interfaces

Advanced users can, if required, configure different replication interfaces to the default for disaster recovery (DR) configurations.

About this task

By default, the eth20 interface is used for data replication in a DR configuration, but you can specify a different interface, if required. You should choose an interface of at least 10 Gb; the performance of a 1 Gb replication link is likely to be poor.

Procedure

- To configure a non-default DR replication link:
 - a) Make the necessary hardware configuration to connect the appliances in the DR appliance.
 - b) Enter the IBM MQ administration mode on both DR appliances by entering the following command:

```
mqcli
```
 - c) Set up the custom interface by entering the following command on both DR appliances:

```
setdrlink -r IP_address
```

Where *IP_address* is the IP address of the interface to be used for replication. You must specify the IP address using ip v4 dotted decimal notation (for example, "192.0.2.8"). The interface that you specify on each appliance must meet the following requirements:

- The IP address specified must map to a valid local interface or aggregation interface (you cannot specify a VLAN interface).
- The local interface must be at least a 10 Gb link, otherwise you are warned.
- You must select the same interface for both appliances in your HA group.

The command sets a new interface from that moment forward, so that any **crtdrprimary** and **crtdrsecondary** commands set up a DR configuration using the newly specified interface. Any queue managers configured to use DR before this point will continue to use the default interface, rather than the new one, unless you remove and recreate the DR configuration for those queue managers.

- To check the configuration of your DR replication link, enter one of the following commands:
 - To view the DR replication link for a particular queue manager:

```
dspdrlink -m QueueManager
```

The IP address and interface name currently configured for that queue manager are displayed, for example:

```
QM1          10.120.5.42 (eth21)
The dspdrlink command succeeded.
```

- To view the default setting, and the settings for existing queue managers:

```
dspdrlink
```

Fore example:

```
Default      10.120.10.100 (eth20)
QMGR         DR Interface (Name)
----         -
QM1          10.120.5.42 (eth21)
QM2          10.120.6.42 (eth22)
The dspdrlink command succeeded.
```

- To revert to the default configuration for the DR replication link:

```
setdrlink -r
```

Related reference

[“setdrlink” on page 685](#)

Specify a non-default link to be used for DR replication in an HA configuration.

[“dspdrlink” on page 683](#)

Display details of the link configured for DR replication in an HA group.

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 288](#).

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 `crtldrprimary` command, specifying an existing queue manager. The queue manager must be stopped when you run this command, unless it is a high availability queue manager (in which case you must leave it to the underlying HA system to handle the stopping of the queue manager).

```
crtldrprimary -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:

```
crtldrprimary -m QM1 -r mydrappl -i 198.51.100.3 -p 2015
```

On successful completion, the command outputs the **crtldrsecondary** 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 `crtldrprimary` command on its successful completion, for example:

```
crtldrsecondary -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 **crtldrsecondary** 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 881](#).

Related reference

[“crtldrprimary” on page 680](#)

Augments an existing queue manager to become the primary queue manager in a disaster recovery configuration.

[“crtldrsecondary” on page 682](#)

Creates a secondary version of a queue manager on the recovery appliance in a disaster recovery configuration.

[“status” on page 881](#)

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.

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 288.

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”](#) on page 332).

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 view the **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 **... > Disaster Recovery** and click **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) Optionally specify a floating IP address. This is used where you are setting up a DR configuration for an HA queue manager. This IP address can be used to replicate data to the DR appliance, regardless of which HA appliance the queue manager is currently running on.
 - i) Click **Create**.
 - j) 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.
 - k) 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 view the **MQ Console**.
 - b) In the console title bar, select **Disaster Recovery > 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:
 - a) Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```
 - b) Ensure that the queue manager is stopped. Enter the following command if necessary:

```
endmqm qmanager
```
 - c) Enter the following command to remove the queue manager from the disaster recovery configuration.

```
dltldrprimary -m qmanager
```
- To remove a secondary queue manager from a disaster recovery configuration:
 - a) Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```
 - b) Enter the following command to remove the queue manager from the disaster recovery configuration and delete it.

```
dltldrsecondary -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:
 - a) Open the console and find the widget that displays the queue manager.
 - b) Ensure that the queue manager is stopped. (Click the stop icon  in the queue manager widget toolbar, if necessary.)

- c) Select ... > **Disaster Recovery** and select **Delete DR Primary**.
- To remove a secondary queue manager from a disaster recovery configuration:
 - a) Open the console and find the widget that displays the queue manager.
 - b) Select ... > **Disaster Recovery** and click **Delete DR Secondary**.

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 an in-progress synchronization has been interrupted. 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 last synchronized. This information is shown only when the status is `Partitioned`, `Inconsistent`, or `Remote appliance(s) unavailable`.
- The time and date that the data was last in sync. This is displayed whenever data is out of sync.

- 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:
 - a) Enter the IBM MQ administration mode by entering the following command:


```
mqcli
```
 - b) 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.
 - c) 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:
 - a) Open the console and find the widget that displays the queue manager.
 - b) Select the queue manager in the widget. Select **More > Disaster Recovery** and click **DR status**.

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 18](#) 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. Do not stop the HA queue manager, the ending of the queue manager is ended by the underlying high availability system. (If you manually end the HA queue manager, the `crtdrprimary` command might hang, in which case contact IBM Support.)
3. Specify that the queue manager is the primary instance in a disaster recovery configuration and optionally include a floating IP address that can be used by either of the appliances in the HA pair:

```
crtdrprimary -m queue_manager_name -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.

V 9.1.5 -f floatingIP

If your HA appliances have static IP addresses assigned to the replication interface (eth20) that are in the same subnet, you can specify a floating IP address. 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. 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.

Note: For versions earlier than IBM MQ Appliance Version 9.1.5 the specification of the floating IP address is mandatory, and appliances in the HA group must be in the same subnet.

In Version 9.1.5, you can currently configure only one HA/DR queue manager with no floating IP address, contact IBM Support for help with overcoming this limitation.

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
```

If you specified a floating IP, then the **-i** parameter is the floating IP, if you omitted the **-f floating_IP** argument, then the **-i** parameter is the static IP address of the replication interface of the appliance that you ran the **crtdrprimary** command on (or the appliance that is the preferred location for the HA queue manager, if that is different). If the queue manager fails over to the other appliance in the HA pair, that appliance identifies itself to the DR appliance so that replication can continue.

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.

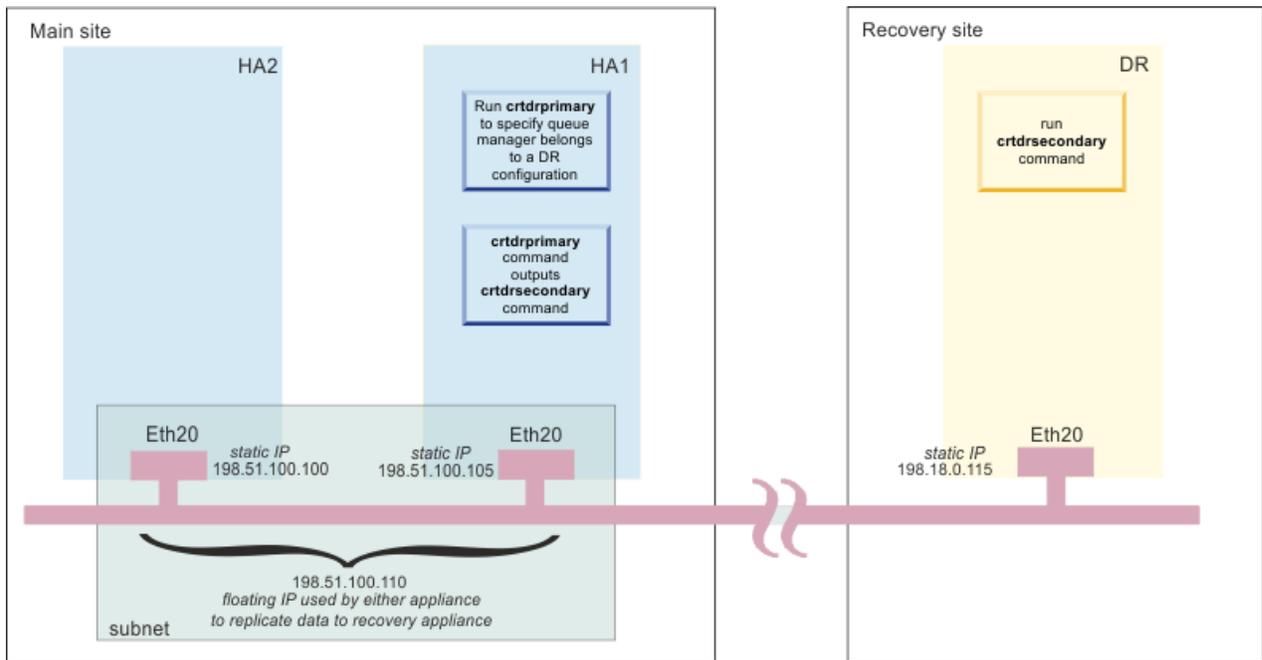


Figure 65. Configuring an HA queue manager for disaster recovery (using example IP addresses in the same subnet)

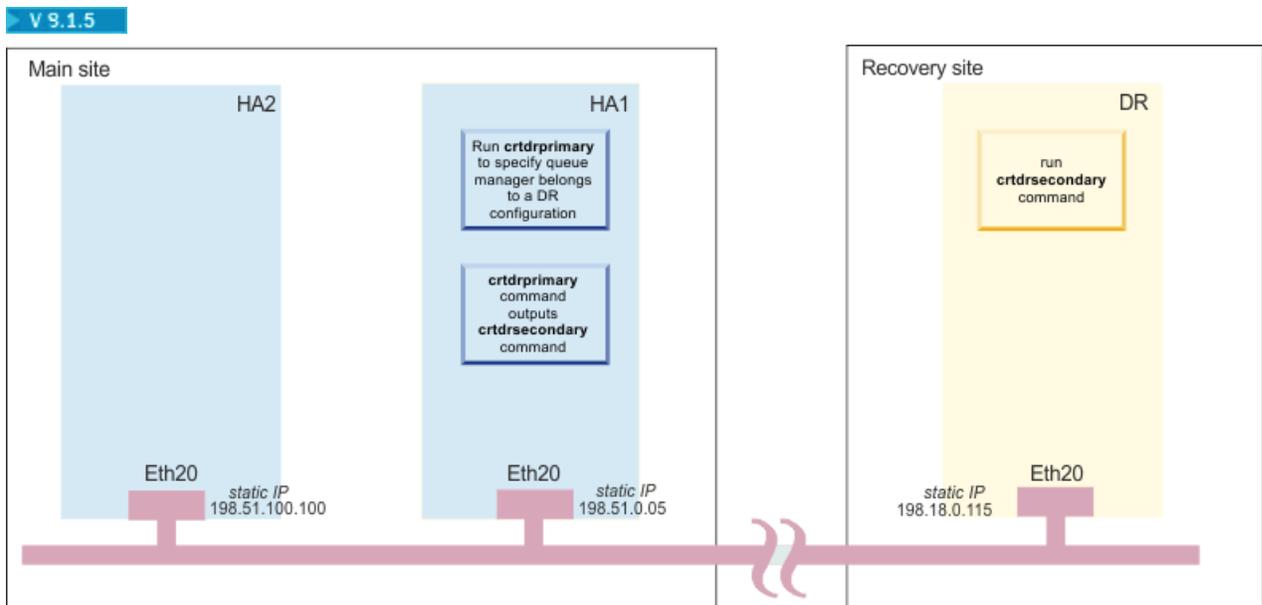


Figure 66. Configuring an HA queue manager for disaster recovery

Related concepts

“Planning network connections” on page 38

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

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.

V 9.1.5 In this example a floating IP is used to replicate data between the HA appliance currently running the queue manager and the DR appliance. The use of a floating IP is optional, and is restricted to the situation where both HA appliances replication (eth20) interfaces are in the same subnet. (For versions earlier than IBM MQ Appliance Version 9.1.5 the specification of the floating IP address is mandatory.)

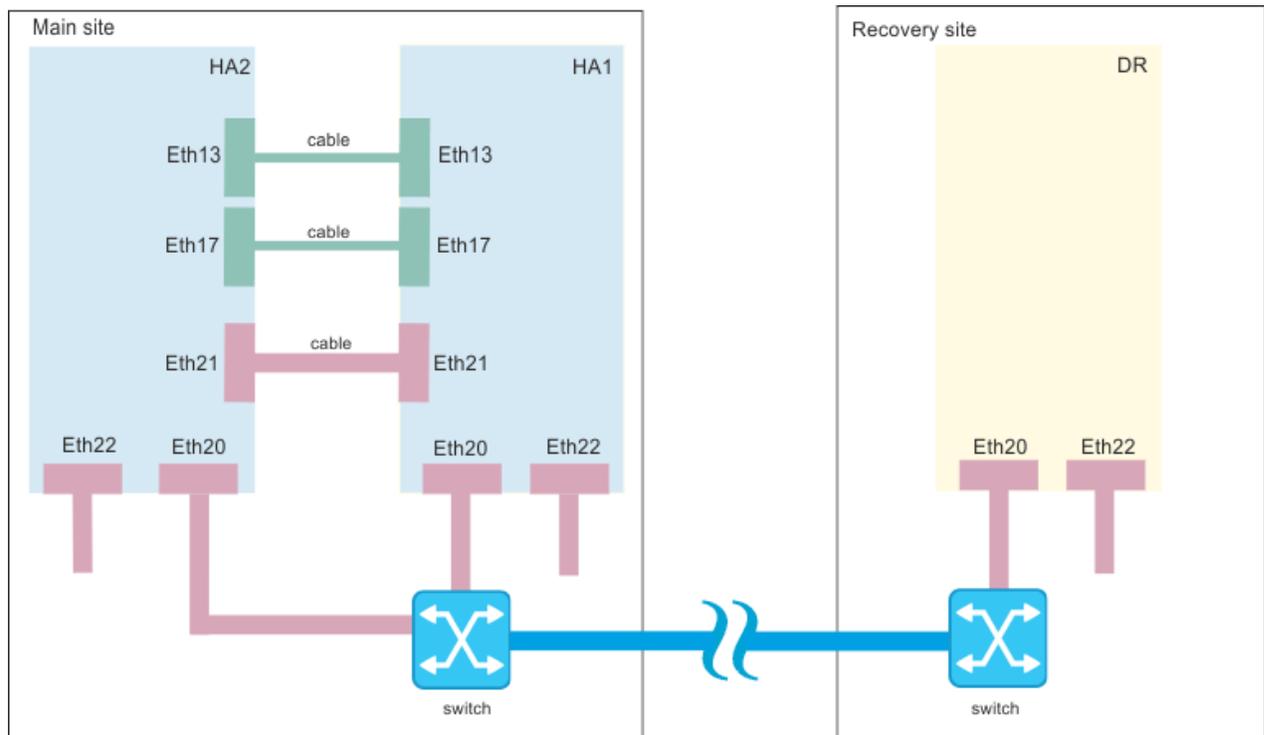


Figure 67. Example HA/DR network configuration

The following tables show how the Ethernet ports are configured on the two HA appliances, HA1 and HA2, and the DR appliance, DR.

Ethernet port	IP address	Description
eth13	192.0.10.11	HA primary group interface
eth17	192.0.20.11	HA group alternative interface
eth20	198.51.100.70	Static IP address configured for DR port
-	198.51.100.19	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

Ethernet port	IP address	Description
eth13	192.0.10.12	HA primary group interface
eth17	192.0.20.12	HA group alternative interface

Table 29. Ethernet ports on appliance HA2 (continued)

Ethernet port	IP address	Description
eth20	198.51.100.12	Static IP address configured for DR port
-	198.51.100.19	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 30. 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

M 2000 M 2001 **Configuring SAN storage**

You can configure the IBM MQ Appliance so that queue managers can use SAN storage.

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances. Support for SAN is now deprecated and there will be no further functional enhancements to this feature.

When you want to use a SAN for queue manager storage you must configure a suitable SAN, configure the appliance to connect to the SAN using the appliance's host bus adapters, and finally create queue managers that use SAN storage.

Queue manager data is not encrypted by the appliance, so you must take steps to secure your SAN storage independently. Volumes configured for appliance use are used to store sensitive information, including certificates and password files relating to the queue manager.

The appliance supports access to a switched SAN fabric accessed by using fibre channel host bus adapters. You must configure your SAN so that each queue manager uses a separate, dedicated LUN.

Ensure that your storage network is zoned so that only the appliance using a particular volume can access the LUN in normal operations. In a disaster recovery scenario, it might be appropriate for multiple appliances to have access to a single LUN, although only one appliance has the volume active (enabled) at any given time.

Ensure administrative access to the data stored on SAN volumes is controlled and audited appropriately.

You must complete the following major steps to configure SAN storage for queue managers on your appliance:

1. Configure the SAN to create the LUNs to be used by queue managers, and allow access from the appliance host bus adapters. LUNs are identified by LUN UUIDs (LUIDs).
2. Define a volume for each LUN, using the LUID to identify it. You require a separate volume for each proposed queue manager, see [“Configuring volumes” on page 302](#).
3. Initialize the file system for the volume. See [“Initializing the file system for a volume” on page 303](#).
4. Create queue managers, specifying the volume for each queue manager, see [“Configuring queue managers to use SAN storage” on page 304](#).

The relationship of SAN and appliance is illustrated in the following figure.

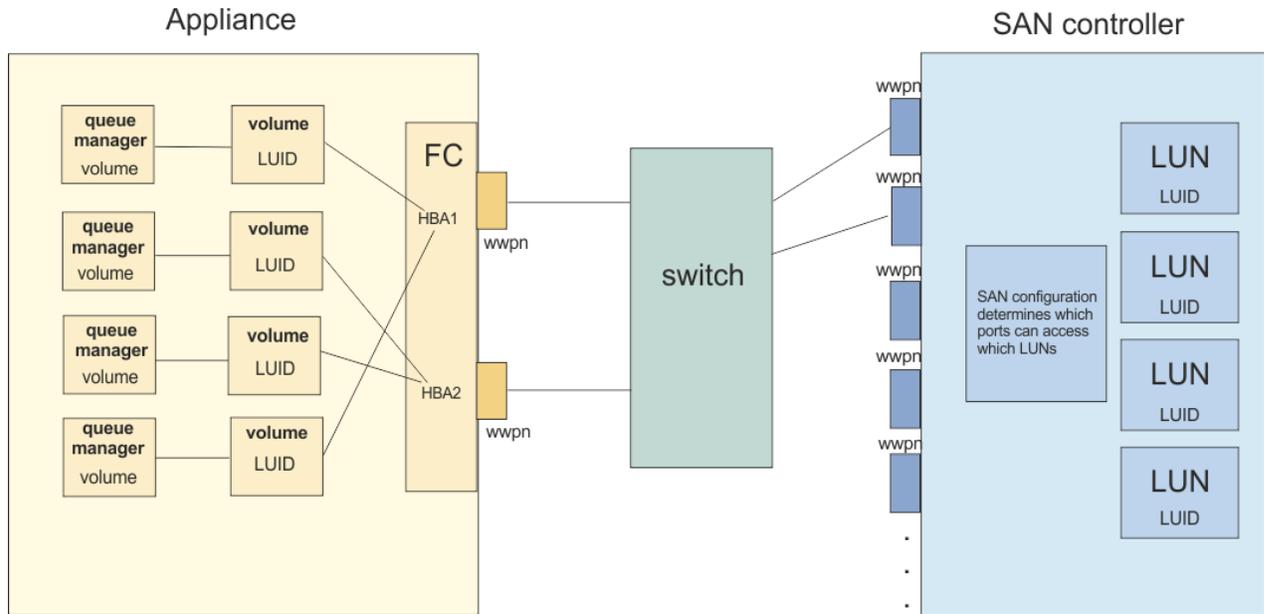


Figure 68. Configuration of SAN and appliance

SAN disks are potentially available to multiple appliances at the same time (one reason for using SAN disks is to obtain higher availability of queue managers). There is a potential risk of more than one appliance trying to use a disk at the same time. To prevent this happening, an IBM MQ Appliance uses SCSI persistent reservations to reserve a disk for its exclusive use. In normal operation, the appliance handles these reservations without involving any administrator action; an appliance reserves the disk before enabling it and then releases the reservation after disabling it. If an appliance attempts to enable a volume that is already reserved by a different appliance, the attempt fails, with appropriate messages reported in the system log. However, if an appliance fails abruptly while holding a reservation on a disk the reservation is not automatically released, and the failed appliance cannot release the reservation either. In this situation, you must remove the SCSI persistent reservation by using the **fibre-channel-unlock-volume** command before another appliance can enable the volume and the queue manager can resume.

► M 2000 M 2001 **Configuring SAN for the appliance**

You must use a storage area network (SAN) that the appliance host bus adapters (HBA) can operate with, and set up SAN to provide storage for queue managers.

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances.

You (or the SAN administrator) must create one SAN device (LUN) for each queue manager that the appliance persists to remote storage. The SAN must be appropriately zoned to permit the appliance or appliances that use a particular LUN to connect to the SAN. The SAN server must report a genuinely unique ID of type 2 or type 3 as defined in the SCSI specification (Vital product data/page 0x80 and 0x83), and this ID must persist across power cycles. The ID (referred to as a LUID or a WWID) is used to identify both the individual LUN and possible routes to that LUN.

The appliance HBA is an Emulex 16Gb FC Dual-port HBA Gen 5 (see “Fibre channel module” on page 78). Full details of the device can be found at <https://www.broadcom.com/products/storage/fibre-channel-host-bus-adapters/lpe16002b>.

M2000 M2001 **Configuring volumes**

You define volume objects on the appliance that represent the storage that a queue manager can access.

Before you begin

Determine the LUID that you will use to identify the LUN used by the queue manager. You can view available LUIDs by choosing **Status > Other Network > Discovered Fibre Channel LUNs** in the IBM MQ Appliance web UI, but your SAN administrator might provide you with the identifiers when the LUNs are created. The LUID is a 64-bit or 128-bit number represented in hexadecimal.

About this task

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances.

You define volume objects that are then used when you create queue managers to define the LUN used by the queue manager.

Procedure

- To configure a volume object by using the IBM MQ Appliance web UI:

- a) Start the IBM MQ Appliance web UI, and click the object icon .
- b) Select **Network Settings > Fibre Channel Volume**.
- c) Click **New**.
- d) Enter a **Name** for the volume.
- e) Click **fibrechannel** to reveal the fibre channel options.
- f) Ensure that **Enable Administrative State** is enabled.
- g) Specify the LUID that identifies the LUN that this volume is used to access.
- h) Select or deselect the **Use Multipath** option as required.
- i) Click **Apply**.

- To configure a volume object by using the command line interface:

- a) Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
- b) Type `config` to enter global configuration mode.
- c) Type the following command to create your volume object and enter volume configuration mode:

```
fibres-channel-volume volume_name
```

- d) Specify the LUID that identifies the LUN that the volume is used to access:

```
lun-uid logical_unit_number
```

- e) Specify whether the volume uses multipath connections or not:

```
use-multipath on | off
```

- f) After you configure the volume object, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Related concepts

[“Fibre channel volume commands” on page 787](#)

You can use the fibre channel volume commands to configure fibre channel volumes on the IBM MQ Appliance.

Related reference

[“fibre-channel-volume” on page 798](#)

This command creates a volume used to connect to a SAN using the appliance fibre channel.

M 2000 M 2001 Initializing the file system for a volume

Before you use a volume for the first time, you must initialize the file system.

About this task

Important: You must only initialize the volume file system once, after you create it and before you attempt to use it. Initializing after you have used the volume erases all the contents of the volume.

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances.

Procedure

- To initialize the file system by using the IBM MQ Appliance web UI
 - a) Start the IBM MQ Appliance web UI, and click the object icon .
 - b) Select **Network Settings > Fibre Channel Volume**.
 - c) Select the volume you are initializing the file space for to view its configuration.
 - d) Disable the administrative state.
 - e) Select **Actions > Initialize File System**
 - f) Enable the administrative state for the volume and save your changes.
- To initialize the file system by using the command line:
 - a) Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
 - b) Type `config` to enter global configuration mode.
 - c) Type the following command to enter volume configuration mode for your volume:

```
fibre-channel-volume volume_name
```

- d) Type the following commands to disable your volume, and leave volume configuration mode:

```
admin-state disabled  
exit
```

- e) Type the following command to initialize the file space for your volume:

```
fibre-channel-fs-init volume_name
```

- f) Type the following commands to re-enter configuration mode for your volume, re-enable your volume, and leave volume configuration mode:

```
fibre-channel-volume volume_name  
admin-state enabled  
exit
```

- g) Type `exit` to leave global configuration mode.

Related reference

[“fibre-channel-fs-init” on page 797](#)

This command initializes the file space for the specified volume.

M 2000 M 2001 Configuring queue managers to use SAN storage

You configure a queue manager to use SAN storage when you create the queue manager.

Before you begin

You must create volume objects before you can create queue managers that use SAN storage. See [“Configuring volumes” on page 302](#).

About this task

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances.

You configure a queue manager to use SAN storage by specifying a volume object when you create the queue manager. The volume object in turn specifies the LUN that the queue manager uses for storage. A queue manager must be uniquely associated with a LUN, you cannot create a queue manager that shares a LUN with another queue manager.

If the appliance on which you are running a queue manager fails, you can re-create the queue manager on another appliance and re-associate it with its SAN storage.

Procedure

- To configure a queue manager by using the IBM MQ Console:
 - a) Open the IBM MQ Console (see [“Using the IBM MQ Console” on page 311](#)).
 - b) In the Queue Manager widget, click the create icon  to start the Create queue manager wizard.
 - c) Enter the name for the queue manager and click **Next**.
 - d) Click **Next** again to skip the high availability set up page (you cannot create a high availability queue manager that uses SAN storage).
 - e) Select **Yes** for **Use external storage**.
 - f) Select the volume object you created for this queue manager from the **SAN volume name** list.
 - g) Select **New** if this is the first time the queue manager volume is used, or **Re-create** if you are recreating a queue manager after an appliance failure, and want to reconnect to SAN storage for that queue manager.
 - h) Click **Create** to create the queue manager.
- To configure a queue manager by using the command line:
 - a) Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as a user with permissions to create IBM MQ objects.
 - b) Enter the `mqcli` command to enter IBM MQ administration mode.
 - c) If you are creating a new queue manager, enter the `crtmqm` command to create a queue manager.

```
crtmqm -fc volume_object qm_name
```

where *volume_object* is the volume object previously created that specifies the LUN that will be exclusively allocated to the queue manager.

You can use other options when you create the queue manager, as described in [“crtmqm \(create queue manager\)” on page 578](#), but you cannot use the `-sx` option to specify high availability capability.

- d) If you are re-creating a queue manager that was running on a failed appliance, enter the `addmqm` command:

```
addmqm -fc volume_object -m qm_name
```

where *volume_object* is the volume object that specifies the LUN that will be exclusively allocated to the queue manager, and *qm_name* is the name of the queue manager that you are recreating.

Related reference

[“crtmqm \(create queue manager\)” on page 578](#)

Create a queue manager.

[“addmqm \(add existing queue manager\)” on page 644](#)

Add an existing queue manager that uses SAN storage.

M 2000 M 2001 Removing queue managers that use SAN storage

You can remove a queue manager that you have configured to use SAN storage.

About this task

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances.

You can remove a SAN queue manager using either the IBM MQ Console or the command line.

You use a different method for removing a queue manager that uses SAN storage than one that uses appliance storage. You use **remove** rather than **delete** (the **rmvmqinf** command rather than the **dltmqm** command), this method leaves the queue manager data intact.

After you have removed the queue manager, you must then remove the volume object used by the queue manager (unless you are intending to reuse it for a different queue manager).

Procedure

- To remove a queue manager by using the IBM MQ Console:
 - a) Open the IBM MQ Console (see [“Using the IBM MQ Console” on page 311](#)).
 - b) Select the queue manager that you want to remove in the queue manager widget and click the delete icon . 
 - c) Click **Remove**. This action removes the instance of the queue manager without removing the queue manager data.
- To remove a queue manager by using the command line:
 - a) Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as a user with permissions to delete IBM MQ objects.
 - b) Enter the `mqcli` command to enter IBM MQ administration mode.
 - c) Enter the following command to remove the queue manager:

```
rmvmqinf qm_name
```

where *qm_name* is the name of the queue manager that you are removing.

What to do next

Remove the volume object that was associated with the removed queue manager. Alternatively, you can associate the volume object with a new queue manager.

Related reference

[“crtmqm \(create queue manager\)” on page 578](#)

Create a queue manager.

[“addmqm \(add existing queue manager\)” on page 644](#)

Add an existing queue manager that uses SAN storage.

[“fibre-channel-volume” on page 798](#)

This command creates a volume used to connect to a SAN using the appliance fibre channel.

Configuring the IBM MQ Console and REST API

The mqweb server that hosts the IBM MQ Console and administrative REST API is provided with a default configuration. You can alter some of this configuration, if required.

Configuring logging for administrative REST API and IBM MQ Console

You can configure the logging levels, maximum log file size, and the maximum number of log files that are used by the mqweb server that hosts the IBM MQ Console and administrative REST API.

Before you begin

You can view the current configuration of the logs by using the **dspmweb properties** command with the **-a** flag. For more information, see [dspmweb](#). You can reset the logging configuration by using the **setmqweb properties** command with the **-k** and **-d** flags. For more information, see [setmqweb](#).

About this task

The log files for the mqweb server can be found in the `:mqtrace/webui` url of the appliance.

Procedure

Use the **setmqweb properties** command from the mqcli prompt to configure logging:

- To set the maximum log file size, use the following command:

```
setmqweb properties -k maxTraceFileSize -v size
```

where *size* specifies the size, in MB, that each log file can reach. The default value is 20.

- To set the maximum number of files to use for logging, use the following command:

```
setmqweb properties -k maxTraceFiles -v max
```

where *max* specifies the maximum number of files. The default value is 2.

- To configure the level of logging that is used, use the following command:

```
setmqweb properties -k traceSpec -v level
```

where *level* is one of the values listed in Table 31 on page 306. The table outlines the logging levels in increasing level of detail. When you enable a logging level, you also enable each level before it. For example, if you enable the ***=warning** logging level, you also enable ***=severe**, and ***=fatal** logging levels.

The default value is ***=info**. Change this value when IBM Service requests it.

Value	Logging level applied
*=off	Logging is turned off.
*=fatal	Task cannot continue and component, application, and server cannot function.
*=severe	Task cannot continue but component, application, and server can still function. This level can also indicate an impending unrecoverable error.

Table 31. Valid logging levels (continued)	
Value	Logging level applied
*=warning	Potential error or impending error. This level can also indicate a progressive failure (for example, the potential leaking of resources).
*=audit	Significant event affecting server state or resources
*=info	General information outlining overall task progress
*=config	Configuration change or status
*=detail	General information detailing subtask progress
*=fine	Trace information - General trace + method entry, exit, and return values
*=finer	Trace information - Detailed trace
*=finest	Trace information - A more detailed trace that includes all the detail that is needed to debug problems
*=all	All events are logged

Configuring the LTPA token expiry interval

When users log in to the IBM MQ Console, an LTPA token is generated. If you use token based authentication with the administrative REST API, an LTPA token is generated when the user logs in using the `/login` REST API resource with the HTTP POST method. The token is used to authenticate the user without the user being required to log in again with their user ID and password, until the token expires. The default expiry interval is 120 minutes, but you can configure when the tokens expire by using the **setmqweb** command.

Before you begin

You can view the current configuration of the token expiry by using the **dspmqweb properties** command with the `-a` flag. For more information, see [dspmqweb](#). You can reset the value of the token expiry by using the **setmqweb properties** command with the `-k` and `-d` flags. For more information, see [setmqweb](#).

Procedure

Use the **setmqweb properties** command from the mqcli prompt to configure the expiry interval:

- Enter the following command:

```
setmqweb properties -k ltpaExpiration -v time
```

where *time* specifies the time, in minutes, before the LTPA token expires and the user is logged out. The default value is 120 minutes.

Configuring the response timeout

By default, the administrative REST API times out if the time taken to send a response back to a client is longer than 30 seconds. You can configure the administrative REST API to use a different timeout value by using the **setmqweb** command.

Before you begin

You can view the current configuration of the timeout by using the **dspmqweb properties** command with the **-a** flag. For more information, see [dspmqweb](#). You can reset the value of the timeout by using the **setmqweb properties** command with the **-k** and **-d** flags. For more information, see [setmqweb](#).

Procedure

Use the **setmqweb properties** command from the mqcli prompt to configure the response timeout:

- Enter the following command:

```
setmqweb properties -k mqRestRequestTimeout -v timeout
```

where *timeout* specifies the time, in seconds, before the time out.

Configuring CORS for the REST API

By default, a web browser does not allow scripts, such as JavaScript, to invoke the administrative REST API when the script is not from the same origin as the REST API. That is, cross-origin requests are not enabled. You can configure Cross Origin Resource Sharing (CORS) to allow cross-origin requests from specified origins.

About this task

You can access the administrative REST API through a web browser, for example through a script. As these requests are from a different origin to the administrative REST API, the web browser refuses the request because it is a cross-origin request. The origin is different if the domain, port, or scheme is not the same.

For example, if you have a script that is hosted at `http://example.com:1999/` you make a cross-origin request if you issue an HTTP GET on a website that is hosted at `https://example.com:9443/`. This request is a cross-origin request because the port numbers and scheme (HTTP) are different.

You can enable cross-origin requests by configuring CORS and specifying the origins that are allowed to access the administrative REST API.

Procedure

Use the **setmqweb properties** command from the mqcli prompt to configure CORS:

- View the current configuration by entering the following command and viewing the `mqRestCorsAllowedOrigins` and `mqRestCorsMaxAgeInSeconds` entries:

```
dspmqweb properties -a
```

- Specify the origins that are allowed to access the administrative REST API by entering the following command:

```
setmqweb properties -k mqRestCorsAllowedOrigins -v allowedOrigins
```

where *allowedOrigins* specifies the origin that you want to allow cross-origin requests from. You can use an asterisk, `*`, to allow all cross-origin requests, or you can enter more than one origin in a comma-separated list. To allow no cross-origin requests, enter empty quotation marks as the value for *allowedOrigins*.

- Specify the time, in seconds, that you want to allow a web browser to cache the results of any CORS pre-flight checks by entering the following command:

```
setmqweb properties -k mqRestCorsMaxAgeInSeconds -v time
```

Example

The following example shows cross-origin requests enabled for `http://example.com:9883`, `https://example.com:1999`, and `https://example.com:9663`. The maximum age of cached results of any CORS pre-flight checks is set to 90 seconds:

```
setmqweb -k mqRestCorsAllowedOrigins -v http://example.com:9883,https://  
example.com:1999,https://example.com:9663  
setmqweb -k mqRestCorsMaxAgeInSeconds -v 90
```


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](#)” on page 577. 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 .

Note: The IBM MQ Console has different timeout behavior to the general IBM MQ Appliance web UI. Because the console can be used for monitoring IBM MQ, the console does not timeout. If you switch to the IBM MQ Appliance web UI, the timeout counter will start. This is set to 600 seconds by default, but can be set to a different value by using the `idle-timeout` command, see “[idle-timeout](#)” on page 997.

Working with queue managers

You can use the queue manager widget in the IBM MQ Console to create, configure, and control local queue managers.

About this task

The queue manager widget lists the queue managers that are running on the appliance. You can select individual queue managers from the list to work with.

You can add a queue manager widget to your dashboard by clicking the add widget icon . Then, select **Local Queue Managers**.

You can rename the widget after you have created it. Hover your mouse pointer in the title bar and the



rename icon appears. Click the icon, enter a new name in the Rename widget window and click **Rename**.

Procedure

- To create a new queue manager, click the create icon  in the local queue manager widget toolbar. The **Create a Queue Manager** wizard opens.
 - a) 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.
 - b) Optional: 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.
 - c) Optionally use **File system size** to specify the size of the file system that is created for the queue manager in either GB or MB.
 - d) Select **Automatic** startup to have the queue manager start automatically when the appliance starts.
 - e) Click **Next** to specify high availability (HA) features if the appliance is part of an HA group, or to continue and specify SAN features (SAN is only available on an M2001 appliance). Otherwise click **Create** to create and start the queue manager.
 - f) Select **Replicated** to specify that the queue manager belongs to a high availability (HA) group.
 - g) Specify the floating IP address that is used to communicate with the queue manager when it is part of an HA group:
 - a. Select the appliance interface that the floating IP address is associated with from the **Floating IP interface** list.
 - b. Specify the floating IP address in IPv4 format in the **Floating IP** field. The address must belong to the same subnet as the static addresses associated with the **Floating IP interface** on each appliance in the group. The subnet mask is already filled in.

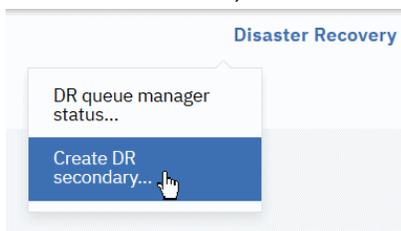
Note: You can select an aggregate interface when specifying the appliance interface that the floating IP address is associated with. If there has been an error in configuring the aggregate interface then it will not appear in the drop-down list. If you are expecting to see an aggregate interface and it does not appear in the list, check the link configuration.
 - h) Click **Next** and select **Yes** to specify SAN features. (SAN features are not available if you have configured high availability, or if you have changed the file system size, or if your appliance is an M2002 model.) Otherwise click **Create** to create and start the queue manager.
 - i) Select the **SAN volume name** for the queue manager (you must configure the volume before you create the queue manager, see [“Configuring volumes”](#) on page 302).
 - j) Click **New** to create a new queue manager, or **Recreate** to re-create a queue manager and attach it to the existing data on the SAN volume.
 - k) Click **Create**. The new queue manager is created and started.
- To start a local queue manager:
 - a) Select the queue manager that you want to start from the list in the local queue manager widget. A toolbar appears.
 - b) Click the start icon .
- To stop a local queue manager:

- a) Select the queue manager that you want to stop from the list in the local queue manager widget. A toolbar appears.
- b) Click the stop icon .
- c) Confirm that you want to stop the queue manager by clicking **Stop**.
- To delete a local queue manager:
 - a) Select the queue manager that you want to delete from the list in the local queue manager widget. A toolbar appears.
 - b) If the queue manager is running, stop it.
 - c) Click the delete icon .
 - d) Confirm that you want to delete the queue manager by clicking **Delete**. The queue manager and all associated objects are deleted. Click **Remove** to remove an instance of the queue manager from SAN storage (M2001 appliances only). This does not delete queue manager data.
- To view and edit the properties of a local queue manager:
 - a) Ensure that the queue manager is running, and select it in the queue manager list.
 - b) Click the properties icon  in the local queue manager widget toolbar. Alternatively, double-click the queue manager.
 - c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line. 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 refresh security for the local queue manager:
 - a) Ensure that the local queue manager is running, and select it in the queue manager list. A toolbar appears.
 - b) Select ... > **Refresh security**
 - c) Select the queue manager security to refresh:
 - Select **Authorization service** to refresh the list of authorizations that is held internally by the authorization services component.
 - Select **Connection authentication** to refresh the cached view of the configuration for connection authentication.
 - Select **SSL** to refresh the cached view of the SSL or TLS key repository. This option also refreshes the locations of the LDAP servers that are used for certified revocation lists, and any cryptographic hardware parameters.
- To work with authority records for the local queue manager:
 - a) Ensure that the local queue manager is running, and select it in the queue manager list. A toolbar appears.
 - b) Select one of the following options:
 - Select ... > **Manage authority records** to work with the authority records for the queue manager, and specify what actions groups of users can take.
 - Select ... > **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 about working with authority records, see [“Working with authority records” on page 329](#).
- To automatically create a dashboard tab for a local queue manager:

- a) Select the queue manager in the local queue manager widget. A toolbar appears.
- b) Select ... > **Add new dashboard tab**
A new dashboard tab is created. The tab has the name of the queue manager.
- To filter the list of local queue managers:
 - a) Type your filter text into the search box.
 - b) To stop filtering, delete the text from the search box.
- To control the high availability operation of a queue manager:
 - a) Select the queue manager in the queue manager list. A toolbar appears.
 - b) Select ... > **High Availability** and select one of the following options:
 - **Add to HA control** - add the selected queue manager to the HA group.
 - **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.
 - **Configure floating IP**. Create or delete a floating IP address for the high availability configuration.
- To add an existing queue manager to a disaster recovery group:
 - a) Select the queue manager and select ... > **Disaster Recovery**. Then select **Create DR Primary**.
 - b) Specify the name of the appliance that hosts the secondary instance of the queue manager (the standby appliance).
 - c) Specify the IP address of the appliance that hosts the secondary instance of the queue manager.
 - d) Specify the port that is used for data replication on both appliances. The port number must be between 1025 - 9999 (do not use port 2222, it is reserved by the appliance).
 - e) Optionally specify a floating IP address. This is used where you are setting up a DR configuration for an HA queue manager. This IP address can be used to replicate data to the DR appliance, regardless of which HA appliance the queue manager is currently running on.
 - f) Click **Create**.
 - g) Copy the command that is displayed in the '**Create DR secondary**' command field when the creation of the DR primary completes. 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:
 - a) In the main toolbar, select **Disaster Recovery** > **Create DR Secondary**



- b) 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 **DR Port** fields with the required values.

- c) Click **Create**.
- To expand the file system of a queue manager:
 - a) Ensure that the queue manager is stopped.
 - b) Select the queue manager, then select ... > **Resize queue manager file system**.
 - c) Specify the new size. This must be the same or greater than the existing size.
 - d) Click **Resize**.

This option is not available for high availability or disaster recovery queue managers, or queue managers that use SAN storage. This option is not available if the queue manager is running.

Note: Resizing file space for a queue manager does involve some I/O, and might degrade the performance of other queue managers while the resize is in progress.

Working with IBM MQ objects

You can use the IBM MQ object widgets in the IBM MQ Console to work with the different types of IBM MQ object.

About this task

Each IBM MQ object widget contains objects that are associated with a specific queue manager. You can add the following types of IBM MQ object widgets to your dashboard:

- Queues widget
- Topics widget
- Listeners widget
- Channels widget
- Client-connection channels widget
- Authentication information widget
- Subscriptions widget

Procedure

- To create an IBM MQ object widget:
 - a) Click **Add widget** .
 - b) Select the appropriate queue manager from the list.
 - c) Click the name of the type of object widget that you want to create.
- To rename the object widget:



- a) Hover your mouse pointer in the title bar and the rename icon appears . Click the icon.
 - b) Enter a new name in the Rename widget window and click **Rename**.
- To configure the IBM MQ object widget:
 - a) Click the configure icon  in the title bar of the widget.
 - b) Optional: Specify the queue manager that the IBM MQ objects are displayed for.
 - c) Optional: Specify whether system objects are shown or hidden.
 - d) Click **Save**.
 - To filter the objects that are displayed in the widget:

- a) Type your filter text into the search box.
- b) To stop filtering, delete the text from the search box.
- To refresh the contents of the widget, click the refresh icon  in the title bar of the widget.
- To remove the widget, click the remove icon  in the title bar of the widget.

Working with queues

You can use the Queues widget in the IBM MQ Console to show the queues that exist for a specific queue manager. You can then add and delete queues, add and clear messages on a queue, browse messages, view and set the properties of a queue, and manage the authority records of a queue.

Before you begin

You must create a queues widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects”](#) on page 315.

About this task

The queues widget lists the queues that exist for a specific queue manager. You can select individual queues from the list to work with.

Procedure

- To add a queue:
 - a) Click the create icon  in the queues widget toolbar.
 - b) Enter a name for the queue. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
 - c) Select the type of queue to add.
 - d) Click **Create**. The new queue is created.
- To put messages to a queue:
 - a) Select the queue that you want to add messages to from the list in the queues widget. You cannot select a model queue.
 - b) Click the put message icon  in the queues widget toolbar.
 - c) Enter the message that you want to put onto the queue.
 - d) Click **Put**.
- To clear messages from a queue:
 - a) Select the local queue that you want to clear messages from the list in the queues widget.
 - b) Select ... > **Clear queue**.
 - c) Confirm that you want to clear the queue by clicking **Clear Queue**.
- To browse messages on a queue:
 - a) Select the local or alias queue that you want to browse from the list in the queues widget.
 - b) Click the browse icon  in the queues widget toolbar. The browse message window opens, displaying messages on the queue.
- To delete a queue:
 - a) Select the queue that you want to delete from the list in the queues widget.
 - b) Click the delete icon  in the queues widget toolbar.

- c) Optional: If the queue has messages on it, confirm that the queue can be cleared by clicking **Clear queue**.
- d) Confirm that you want to delete the queue by clicking **Delete**. The queue is deleted.
- To view and edit the properties of a queue:
 - a) Select the queue in the queues widget.
 - b) Click the properties icon  in the queues widget toolbar. Alternatively, double-click the queue.
 - c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line. For information about a property, view the property information in [IBM Knowledge Center](#).
- To view and edit authority records for a queue:
 - a) Select the queue in the widget.
 - b) Click ... > **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” on page 329](#).

Working with topics

You can use the topics widget in the IBM MQ Console to add and delete topics, and view and set the properties of a topic.

Before you begin

You must create a topics widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

About this task

The topics widget lists the topics that exist for a specific queue manager. You can select individual topics from the list to work with.

Procedure

- To add a topic:
 - a) Click the create icon  in the topics widget toolbar.
 - b) Enter a name for the new topic. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
 - c) Specify the topic string that you publish messages for the topic to. For more information, see [IBM Knowledge Center](#).
 - d) Click **Create**. The new topic is created.
- To delete a topic:
 - a) Select the topic that you want to delete from the list in the topics widget.
 - b) Click the delete icon  in the topics widget toolbar.
 - c) Confirm that you want to delete the topic by clicking **Delete**. The topic is deleted.
- To view and edit the properties of a topic:
 - a) Select the topic in the topics widget.
 - b) Click the properties icon  in the topics widget toolbar. Alternatively, double-click the topic.

- c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line. For information about a property, view the property information in [IBM Knowledge Center](#).
- To publish a message on a topic:
 - a) Click the put message icon  in the topics widget toolbar.
 - b) Enter a message in the **Message** field.
 - c) Enter the topic string to publish the message on in the **Topic string** field.
 - d) Click **Publish**.
- To subscribe to a topic:
 - a) Click the subscribe icon  in the topics widget toolbar.
 - b) Enter the topic string to subscribe to in the **Topic string** field.
 - c) Click **Subscribe**.
- To view and edit authority records for a topic:
 - a) Select the topic in the topics widget.
 - b) Click the ... > **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” on page 329](#).

Working with listeners

You can use the listeners widget in the IBM MQ Console to add and delete listeners, start and stop listeners, view and set listener properties, and manage the authority records for a listener.

Before you begin

You must create a listeners widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

About this task

The listeners widget lists the listeners that exist for a specific queue manager. You can select individual listeners from the list to work with.

Procedure

- To add a TCP/IP listener:
 - a) Click the create icon  in the listeners widget toolbar.
 - b) Enter a name for the listener. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
 - c) Enter an available TCP/IP port for the listener. The port number must not exceed 65535.
 - d) Click **Create**. The new listener is created.
- To delete a listener:
 - a) Select the listener that you want to delete from the list in the listeners widget.
 - b) Click the delete icon  in the listeners widget toolbar.
 - c) Confirm that you want to delete the listener by clicking **Delete**. The listener is deleted.
- To start a listener:
 - a) Select the listener that you want to start from the list in the listeners widget.

- b) Click the start icon  in the listeners widget toolbar.
 - To stop a listener:
 - a) Select the listener that you want to stop from the list in the listeners widget.
 - b) Click the stop icon  in the listeners widget toolbar.
 - c) Confirm that you want to stop the listener by clicking **Stop**.
 - To view and edit the properties of a listener:
 - a) Select the listener in the listeners widget.
 - b) Click the properties icon  in the listeners widget toolbar. Alternatively, double-click the listener.
 - c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line. For information about a property, view the property information in [IBM Knowledge Center](#).
- Note:** By default, listeners are created with the **control** property set to **manual**. 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:
 - a) Select the listener in the listeners widget.
 - b) Click ... > **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” on page 329](#).

Working with channels

You can use the channels widget in the IBM MQ Console to add and delete channels, start and stop channels, reset and resolve channels, and ping channels. You can also view and set the properties of a channel, and manage authority records for the channel.

Before you begin

You must create a channels widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

About this task

The channels widget lists the channels that exist for a specific queue manager. You can select individual channels from the list to work with.

Procedure

- To add a channel:
 - a) Click the create icon  in the channels widget toolbar.
 - b) Enter a name for the channel. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
 - c) Select the type of channel to add.
 - d) 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)*. *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.
 - e) 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.

- f) Click **Create**. The new channel is created.
- To delete a channel:
 - a) Select the channel that you want to delete from the list in the channels widget.
 - b) Click the delete icon  in the widget toolbar.
 - c) Confirm that you want to delete the channel by clicking **Delete**. The channel is deleted.
- To start a channel:
 - a) Select the channel that you want to start from the list in the channels widget.
 - b) Click the start icon  in the widget toolbar.
- To stop a channel:
 - a) Select the channel that you want to stop from the list in the channels widget.
 - b) Click the stop icon  in the widget toolbar.
 - c) Confirm that you want to stop the channel by clicking **Stop**.
- To view the properties of a channel:
 - a) Select the channel in the channels widget.
 - b) Click the properties icon  in the channels widget toolbar. Alternatively, double-click the channel.
 - c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line. For information about a property, view the property information in [IBM Knowledge Center](#).
- To reset a channel:
 - a) Select the channel in the channels widget.
 - b) Click ... > **Reset**.
 - c) 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.
 - d) Click **Reset Channel**.
- To resolve a channel:
 - a) Select the channel in the channels widget.
 - b) Click ... > **Resolve**.
 - c) Choose whether to commit or back out the current batch of messages by clicking **Commit** or **Back out**.
- To ping a channel:
 - a) Select the channel in the channels widget.
 - b) Click ... > **Ping**.
- To view or edit authority records for a channel:
 - a) Select the channel in the widget.
 - b) Click ... > **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”](#) on page 329.

Working with client-connection channels

You can use the client-connection channels widget in the IBM MQ Console to add and delete client-connection channels on a queue manager, view and set the properties, and manage the authority records for the channel.

Before you begin

You must create a client-connection channels widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects”](#) on page 315.

About this task

The client-connection channels widget lists the client-connection channels that exist for a specific queue manager. You can select individual client-connection channels from the list to work with.

Procedure

- To add a client-connection channel:
 - a) Click the create icon  in the client-connection channels widget toolbar.
 - b) Enter a name for the client-connection channel. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
 - c) 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.
 - d) Click **Create**. The new client-connection channel is created.
- To delete a client-connection channel:
 - a) Select the client-connection channel that you want to delete from the list in the client-connection channels widget.
 - b) Click the delete icon  in the widget toolbar.
 - c) 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:
 - a) Select the client-connection channel in the client-connection channels widget.
 - b) Click the properties icon  in the client-connection channels widget toolbar. Alternatively, double-click the client-connection channel.
 - c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line. For information about a property, view the property information in [IBM Knowledge Center](#).
- To view and edit authority records for a client-connection channel:
 - a) Select the client-connection channel in the client-connection channels widget.
 - b) Click **... > 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”](#) on page 329.

Working with authentication information

You can use the authentication information widget in the IBM MQ Console to add and delete authentication information objects on a queue manager. You can also view and set the properties, and manage the authority records for the objects.

Before you begin

You must create an authentication information widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects”](#) on page 315.

About this task

The authentication information widget lists the authentication information that exists for a specific queue manager. You can select individual authentication information from the list to work with.

The queue manager authentication information forms part of IBM MQ support for Transport Layer Security (TLS). These objects contain the definitions that are required to perform certificate revocation checking by using OCSP or Certificate Revocation Lists (CRLs) on LDAP servers, and the definitions that are required to enable user ID and password checking.

Procedure

- To add an authentication information object:

- a) Click the create icon  in the authentication information widget toolbar.
- b) Specify the name of the authentication information object. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
- c) Specify the type of authentication information object.
- d) Specify additional information appropriate to the object type:
 - For **CRL LDAP**, specify the **LDAP server name**. This name 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.
 - For **OCSP**, specify the **OCSP responder URL**. This URL is the URL of the responder that is 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.
 - For **IDPW OS**, there are no additional requirements.
 - For **IDPW LDAP**, specify the **LDAP server name** and the **Short user** name. The LDAP server name 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. The short user name is the field in the LDAP user record that is used as a short name for the connection.
- e) Click **Create**.

- To delete an authentication information object:

- a) Select the authentication information object that you want to delete from the list in the widget.
- b) Click the delete icon  in the widget toolbar.
- c) 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:

- a) Select the authentication information object in the widget.
- b) Click the properties icon  in the widget toolbar. Alternatively, double-click the authentication information object.

- c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line.
- To view and edit authority records for an authentication information object:
 - a) Select the authentication information object in the authentication information widget.
 - b) Click ... > **Manage Authority Records**. The authority records show the permissions that users and administrators have on the selected authentication information object. For details of editing the authority records, see [“Working with authority records” on page 329](#).

Working with subscriptions

You can use the subscriptions widget in the IBM MQ Console to add and delete subscriptions on a queue manager, view and set the properties, and manage the authority records for the subscriptions.

Before you begin

You must create a subscriptions widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

About this task

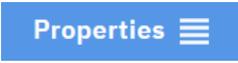
Subscriptions are issued to a queue manager and contain information about the publications the subscriber wants to receive:

- The topic string that the subscriber is interested in; this topic can resolve to multiple topic strings if wildcards are used.
- An optional selection string to be applied to published messages.
- The name of the queue on which selected publications are placed.

For more information about subscriptions, see [Subscribers and subscriptions](#) and [DEFINE SUB](#) in the IBM MQ documentation.

Procedure

- To add a subscription object:
 - a) Click the create icon  in the subscriptions widget toolbar.
 - b) Specify the name of the object. Valid characters are letters and numbers and the ".", "/", "_", and "%" characters.
 - c) Select a **Destination class** of **Managed** or **Provided**. If you select **Managed**, a destination is created on the local queue manager.
 - d) 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.
 - e) In the **Topic string** field, specify the topic string to subscribe to.
 - f) 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.
 - g) 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 that are published on the topic only within this queue manager.
 - h) Optional: Specify a **Selector**. A selection string is an expression that is applied to a publication to determine whether it matches a subscription.
 - i) Click **Create**.
- To delete a subscription object:
 - a) Select the subscription object that you want to delete from the list in the subscriptions widget.

- b) Click the delete icon  in the widget toolbar.
- c) 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:
 - a) Select the subscription object in the widget.
 - b) Click the properties icon  in the widget toolbar. Alternatively, double-click the subscription object.
 - c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line.

Working with channel authentication records

You can use the channel authentication records widget in the IBM MQ Console to add and delete channel authentication records on a queue manager. You can also view and set the properties for channel authentication records.

Before you begin

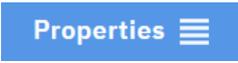
You must create a channel authentication records widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects”](#) on page 315.

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 product documentation.

Procedure

- To add a channel authentication record with an SSL/TLS distinguished name identity, see [“Creating channel authentication records with an SSL/TLS Distinguished Name identity”](#) on page 325.
- To add a channel authentication record with a client application user ID identity, see [“Creating channel authentication records with a client application user ID identity”](#) on page 326.
- To add a channel authentication record with a remote queue manager name identity, see [“Creating channel authentication records with a remote queue manager name identity”](#) on page 327.
- To add a channel authentication record with an address identity, see [“Creating channel authentication records with an IP address identity”](#) on page 327.
- To delete a channel authentication record:
 - a) Select the channel authentication record that you want to delete from the list in the channel authentication records widget.
 - b) Click the delete icon  in the widget toolbar.
 - c) 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:
 - a) Select the channel authentication record that you want to edit from the list in the channel authentication record widget.
 - b) Click the properties icon  in the widget toolbar. Alternatively, double-click the channel authentication record.

- c) View the properties and edit them as required. If the property text box is disabled, the property is read-only, or can be edited only from the command line.

Creating channel authentication records with an SSL/TLS Distinguished Name identity

You can use the channel authentication records widget to create allowing, blocking, and warning channel authentication records with an SSL/TLS Distinguished Name identity. The SSL/TLS distinguished name identity matches to users who present an SSL or TLS personal certificate that contains a specified Distinguished Name.

Before you begin

You must create a channel authentication records widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

Procedure

- To add a channel authentication record:

- a) Click the create icon  in the channel authentication record widget toolbar.
- b) Select the **Rule Type** to indicate what type of rule you want on the channel authentication record:
 - Select **Allow** to allow access to inbound connections.
 - Select **Block** to block access to inbound connections.
 - Select **Warn** to warn about access to inbound connections that would be blocked. The connection is allowed access, and an error message is reported. If events are configured, an event message is created that shows the details of what would be blocked. Only matched rules are reported.
- c) Select the **SSL/TLS Distinguished Name** identity type from the list.
- d) Click **Next**
- e) Specify a **Channel profile**.

The channel profile is the name of the channel or set of channels for which you are setting the channel authentication. The profile can contain wildcards so that you can block a range of channels. For example, the profile `alphadelta*` blocks channels named `alphadelta1`, `alphadelta2`, `alphadelta3` and so on.
- f) Specify the **Peer name**. For example, `CN=John Smith, O=IBM ,OU=Test , C=GB`. For more information about peer names, see [WebSphere MQ rules for SSLPEER values](#) in the IBM MQ product documentation.
- g) Optional: Specify the **Address** filter that is used. The address is the IP address that is expected at the other end of the channel.
- h) Optional: Specify the **SSL cert issuer name**. The SSL cert issuer name is the name of the certificate authority that the SSL/TLS certificate must be issued by.
- i) Optional: Click **Next**.
- j) Optional: For an **Allow** rule type, you can optionally specify the **User source** for the channel authentication record. The user source specifies the source of the user ID that is used when the inbound connection matches the SSL/TLS Distinguished Name.
 - The **Channel** option specifies that inbound connections that match the mapping use the flowed user ID or any user that is defined on the channel object.
 - The **Map** option specifies that inbound connections that match the mapping use the user ID that is specified in the **MCA user ID** field.
- k) Optional: Click **Next**.
- l) Optional: Specify a **Description** for the channel authentication record.
- m) Click **Create**. The new channel authentication record is created.

Creating channel authentication records with a client application user ID identity

You can use the channel authentication records widget to create allowing, blocking, and warning channel authentication records with a client application user ID identity. The client application user ID identity matches to client application IDs from a client-connection channel.

Before you begin

You must create a channel authentication records widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

Procedure

- To add a channel authentication record:

- Click the create icon  in the channel authentication record widget toolbar.
- Select the **Rule Type** to indicate what type of rule you want on the channel authentication record:
 - Select **Allow** to allow access to inbound connections.
 - Select **Block** to block access to inbound connections.
 - Select **Warn** to warn about access to inbound connections that would be blocked. The connection is allowed access, and an error message is reported. If events are configured, an event message is created that shows the details of what would be blocked. Only matched rules are reported.
- Select the **Client application user ID** identity type from the list.
- Click **Next**.
- Specify a **Channel profile**.

The channel profile name is the name of the channel or set of channels for which you are setting the channel authentication. The profile can contain wildcards so that you can block a range of channels. For example, the profile `a1phade1ta*` blocks channels named `alphadelta1`, `alphadelta2`, `alphdelta3` and so on.
- Specify the **Client user ID**. The client user ID is the user ID of the client that you want to allow, block, or warn about.
- Optional: Specify the **Address** filter that is used. The address is the IP address that is expected at the other end of the channel.
- Optional: Click **Next**.
- Optional: For an **Allow** rule type, you can optionally specify the **User source** for the channel authentication record. The user source specifies the source of the user ID that is used when the inbound connection matches the client user ID.
 - The **Channel** option specifies that inbound connections that match the mapping use the flowed user ID or any user that is defined on the channel object.
 - The **Map** option specifies that inbound connections that match the mapping use the user ID that is specified in the **MCA user ID** field.
- Optional: Click **Next**.
- Optional: Specify a **Description** for the channel authentication record.
- Click **Create**. The new channel authentication record is created.

Creating channel authentication records with a remote queue manager name identity

You can use the channel authentication records widget to create allowing, blocking, and warning channel authentication records with a remote queue manager name identity. The remote queue manager name identity matches to the specified queue manager.

Before you begin

You must create a channel authentication records widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

Procedure

- To add a channel authentication record:
 - a) Click the create icon  in the channel authentication record widget toolbar.
 - b) Select the **Rule Type** to indicate what type of rule you want on the channel authentication record:
 - Select **Allow** to allow access to inbound connections.
 - Select **Block** to block access to inbound connections.
 - Select **Warn** to warn about access to inbound connections that would be blocked. The connection is allowed access, and an error message is reported. If events are configured, an event message is created that shows the details of what would be blocked. Only matched rules are reported.
 - c) Select the **Remote queue manager name** identity type from the list.
 - d) Click **Next**
 - e) Specify a **Profile Name**.

The profile name is the name of the channel or set of channels for which you are setting the channel authentication. The profile can contain wildcards so that you can block a range of channels. For example, the profile `alphadelta*` blocks channels named `alphadelta1`, `alphadelta2`, `alphadelta3` and so on.
 - f) Specify the **Queue manager name**. The queue manager name specifies the name of the remote queue manager that you want to allow, block, or warn about.
 - g) Optional: Specify the **Address** filter that is used. The address is the IP address that is expected at the other end of the channel.
 - h) Optional: Click **Next**.
 - i) Optional: For an **Allow** rule type, you can optionally specify the **User source** for the channel authentication record. The user source specifies the source of the user ID that is used when the inbound connection matches the remote queue manager name.
 - The **Channel** option specifies that inbound connections that match the mapping use the flowed user ID or any user that is defined on the channel object.
 - The **Map** option specifies that inbound connections that match the mapping use the user ID that is specified in the **MCA user ID** field.
 - j) Optional: Click **Next**.
 - k) Optional: Specify a **Description** for the channel authentication record.
 - l) Click **Create**. The new channel authentication record is created.

Creating channel authentication records with an IP address identity

You can use the channel authentication records widget to create allowing, blocking, and warning channel authentication records with an address identity. The address identity matches to specific IP addresses.

Before you begin

You must create a channel authentication records widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

Procedure

- To add a channel authentication record:
 - a) Click the create icon  in the channel authentication record widget toolbar.
 - b) Select the **Rule Type** to indicate what type of rule you want on the channel authentication record:
 - Select **Allow** to allow access to inbound connections.
 - Select **Block** to block access to inbound connections.
 - Select **Warn** to warn about access to inbound connections that would be blocked. The connection is allowed access, and an error message is reported. If events are configured, an event message is created that shows the details of what would be blocked. Only matched rules are reported.
 - c) Select the **Address** identity type from the list.
 - d) Click **Next**
 - e) Optional: For a **Block** or **Warn** rule type, specify **When to match**.
You can choose from these options:
 - **At the listener**. This option attempts to match the rule at the listener.
 - **At the channel**. This option attempts to match the rule at the channel.
 - f) Specify a **Profile Name**.
The profile name is the name of the channel or set of channels for which you are setting the channel authentication. The profile can contain wildcards so that you can block a range of channels. For example, the profile `alphadelta*` blocks channels named `alphadelta1`, `alphadelta2`, `alphadelta3` and so on.
 - g) Specify an **Address**. The address is the IP address or a comma-separated list of IP addresses that are allowed or blocked.
 - h) Optional: Click **Next**.
 - i) Optional: For an **Allow** rule type, you can optionally specify the **User source** for the channel authentication record. The user source specifies the source of the user ID that is used when the inbound connection matches the remote queue manager name.
 - The **Channel** option specifies that inbound connections that match the mapping use the flowed user ID or any user that is defined on the channel object.
 - The **Map** option specifies that inbound connections that match the mapping use the user ID that is specified in the **MCA user ID** field.
 - j) Optional: Click **Next**.
 - k) Optional: Specify a **Description** for the channel authentication record.
 - l) Click **Create**. The new channel authentication record is created.

Creating channel authentication records with a final assigned user ID identity

You can use the channel authentication records widget to create blocking and warning channel authentication records with a final assigned user ID identity. The final assigned user ID identity matches to list of specified user IDs from a server channel.

Before you begin

You must create a channel authentication records widget before you can use it. For more information about creating IBM MQ object widgets, see [“Working with IBM MQ objects” on page 315](#).

Procedure

- To add a channel authentication record:

- a) Click the create icon  in the channel authentication record widget toolbar.
- b) Select the **Rule Type** to indicate what type of rule you want on the channel authentication record:
- Select **Block** to block access to inbound connections.
 - Select **Warn** to warn about access to inbound connections that would be blocked. The connection is allowed access, and an error message is reported. If events are configured, an event message is created that shows the details of what would be blocked. Only matched rules are reported.
- c) Select the **Final assigned user ID** identity type from the list.
- d) Click **Next**
- e) Specify a **Profile Name**.
- The profile name is the name of the channel or set of channels for which you are setting the channel authentication. The profile can contain wildcards so that you can block a range of channels. For example, the profile `alphadelta*` blocks channels named `alphadelta1`, `alphadelta2`, `alphadelta3` and so on.
- f) Specify the **User list**. The user list is a comma-separated list of user IDs to be blocked from the channel.
- g) Optional: Click **Next**.
- h) Optional: Specify a **Description** for the channel authentication record.
- i) Click **Create**. The new channel authentication record is created.

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

You can fine-tune the access that a group of messaging users has to a particular queue manager or IBM MQ object by using authority records. 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 'qm3'

Delete 		1 item selected	Cancel
▲ Entity name	Entity type		
mqm	Group		
mqsystem	User		
Total: 2		Last updated: 3:40:18 PM	

Administration

- Change
- Delete
- Display
- Ctrl

Context

- Set all context
- Set identity context

MQI

- Alternate user authority
- Connect
- Inquire
- Set
- System

Check all

Uncheck all

Close

Save

Authority records for 'q1' on qm3

Delete 		1 item selected	Cancel
▲ Entity name	Entity type		
mqm	Group		
mqsystem	User		

Administration	Context	MQI
<input checked="" type="checkbox"/> Change	<input checked="" type="checkbox"/> Pass all context	<input checked="" type="checkbox"/> Browse
<input checked="" type="checkbox"/> Clear	<input checked="" type="checkbox"/> Pass identity context	<input checked="" type="checkbox"/> Inquire
<input checked="" type="checkbox"/> Delete	<input checked="" type="checkbox"/> Set all context	<input checked="" type="checkbox"/> Get
<input checked="" type="checkbox"/> Display	<input checked="" type="checkbox"/> Set identity context	<input checked="" type="checkbox"/> Put
		<input checked="" type="checkbox"/> Set

Procedure

- To view or edit an authority record for an IBM MQ object:
 - a) Select the object in a widget on the dashboard. The associated queue manager must be running.
 - b) From the appropriate widget toolbar, select ... > **Manage Authority Records**.
 - c) Select the group that you want to view the authority record for. The authorities for that group are displayed.
 - d) Select or clear authorities as required. Different authorities are available depending on the type of object that you are creating an authority record for.
 - e) Click **Save**.
- To view or edit a create authority record for a queue manager:
 - a) Select the queue manager in a queue manager widget on the dashboard. The queue manager must be running.
 - b) From the widget toolbar, select ... > **Manage Create Authority Records**.
 - c) Select the group that you want to view the create authority record for. The authorities for that group are displayed.
 - d) Select or clear create authorities as required.
 - e) Click **Save**.
- To create an authority record for an IBM MQ object:

- a) Select the IBM MQ object in a widget on the dashboard. The associated queue manager must be running.
 - b) From the widget toolbar, select ... > **Manage Authority Records**.
 - c) Click the plus icon .
 - d) Specify the name of the user or group that you are creating the authority record for. The user or group must exist.
 - e) Select the **Entity Type** to specify whether the entity is a user or a group.
 - f) Click **Create**.
 - g) Select or clear the authorities that you want the user or group to have. Different authorities are available for each type of object.
 - h) Click **Save**.
- To create an authority record for creating objects on a queue manager:
 - a) Select the queue manager in a widget on the dashboard. The queue manager must be running.
 - b) From the widget toolbar, select ... > **Manage Create Authority Records**.
 - c) Click the create icon .
 - d) Specify the name of the user or group that you are creating the authority record for. The user or group must exist.
 - e) Select the **Entity Type** to specify whether the entity is a user or a group.
 - f) Click **Create**.
 - g) Select or clear the create authorities that you want the user or group to have.
 - h) Click **Save**.

Monitoring system resource usage

You use the Charts widget in the IBM MQ Console to view monitoring data for queue managers.

About this task

You add a Charts widget to your dashboard and then configure it to monitor a particular aspect of resource usage. You can create many instances of the Charts widget to display different data. The data is displayed in a chart format.

In order to start monitoring system resource, the user or group must be part of MQWebAdmin or MQWebAdminRO. Once the monitoring is started, any authenticated MQWebUser can access the stream and monitor by adding the Charts widget.

Data is collected at 10-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 must have at least one running queue manager before you can configure a chart widget.

Procedure

1. Add a Charts widget to your dashboard:

- a) Click the Add widget icon .
- b) Select **Charts**.

2. Configure the Charts widget to show data:

- a) Click the configure icon  in the title bar of the Charts widget.

b) Optional: Enter a **Widget title**. This title is shown in the title bar of the widget.

c) Select the **Resource class** to monitor:

Platform central processing units

Monitor the usage of the CPUs.

Platform persistent data stores

Monitor the use of disk resource.

API usage statistics

Monitor API calls.

API per-queue usage statistics

Monitor API calls by individual queues. When you choose this class, you specify the queue name to monitor in the **Object** field.

d) Select the **Resource type** to monitor.

The resource types that are available to select depend on the resource class that is selected. The following table shows the 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.
	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.
	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 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.

<i>Table 32. Resource types (continued)</i>		
Class	Type	Description
	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 sync points by the queue manager.
	Subscribe	Select this type to view data that is 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.

e) Select the **Resource element** to monitor:

The resource elements that are available to select depend on the resource class and resource type that are selected. The following tables show the resource elements:

<i>Table 33. Elements for Platform central processing units resources</i>		
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 load average over 1 minute.
	CPU load - five-minute average	Shows the load average over 5 minutes.
	CPU load - fifteen-minute average	Shows the load average over fifteen minutes.
	RAM free percentage	Shows the percentage of free RAM memory.

<i>Table 33. Elements for Platform central processing units resources (continued)</i>		
Type	Element	Description
	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 that are 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 that are 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.

<i>Table 34. Elements for Platform persistent data stores resources</i>		
Type	Element	Description
Disk usage - platform wide	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 free.
	MQ FDC file count	Shows the current number of FDC files.
	Appliance data - bytes in use	Shows the overall disk usage.
	Appliance data - free space	Shows the overall free space.
	System volume - bytes in use	
	System volume - free space	
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 that is reserved for queue manager files that is free.

<i>Table 34. Elements for Platform persistent data stores resources (continued)</i>		
Type	Element	Description
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 is 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.

<i>Table 35. Elements for API usage statistics resources</i>		
Type	Element	Description
MQCONN and MQDISC	MQCONN/MQCONN count	Shows the number of calls to MQCONN and MQCONN.
	Failed MQCONN/MQCONN count	Shows the number of failed calls to MQCONN and MQCONN.
	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.

<i>Table 35. Elements for API usage statistics resources (continued)</i>			
Type	Element	Description	
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 that are removed from queues by MQGET.	

<i>Table 35. Elements for API usage statistics resources (continued)</i>		
Type	Element	Description
	Interval total destructive get - byte count	Bytes of data that is removed from queues by MQGET.
	Non-persistent message destructive get - count	Number of non-persistent messages that are removed from queues by MQGET.
	Persistent message destructive get - count	Number of persistent messages that are 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 that are returned to MQGET.
	Got persistent messages - byte count	Shows a count of bytes of persistent messages that are 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.

<i>Table 35. Elements for API usage statistics resources (continued)</i>		
Type	Element	Description
	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

<i>Table 35. Elements for API usage statistics resources (continued)</i>		
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 that are 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 that are published to subscribers.
	Published to subscribers - byte count	Shows the byte count of messages that are published to subscribers.
	Non-persistent - topic MQPUT/MQPUT1 count	Shows the number of non-persistent messages that are put to topics.
	Persistent - topic MQPUT/MQPUT1 count	Shows the number of persistent messages that are put to topics.
	Failed topic MQPUT/MQPUT1 count	Shows the number of failed attempts to put to a topic.

<i>Table 36. Elements for API per-queue usage statistics resources</i>		
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.

<i>Table 36. Elements for API per-queue usage statistics resources (continued)</i>		
Type	Element	Description
	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	
	Queue avoided bytes	
	Lock contention	
MQGET	MQGET count	
	MQGET byte count	
	Destructive MQGET non-persistent message count	Number of non-persistent messages that are removed from the queue by MQGET.
	Destructive MQGET persistent message count	Number of persistent messages that are removed from the queue by MQGET.
	Destructive MQGET non-persistent byte count	Shows a count of bytes of non-persistent messages that are returned to MQGET.
	Destructive MQGET persistent byte count	Shows a count of bytes of persistent messages that are returned to MQGET.

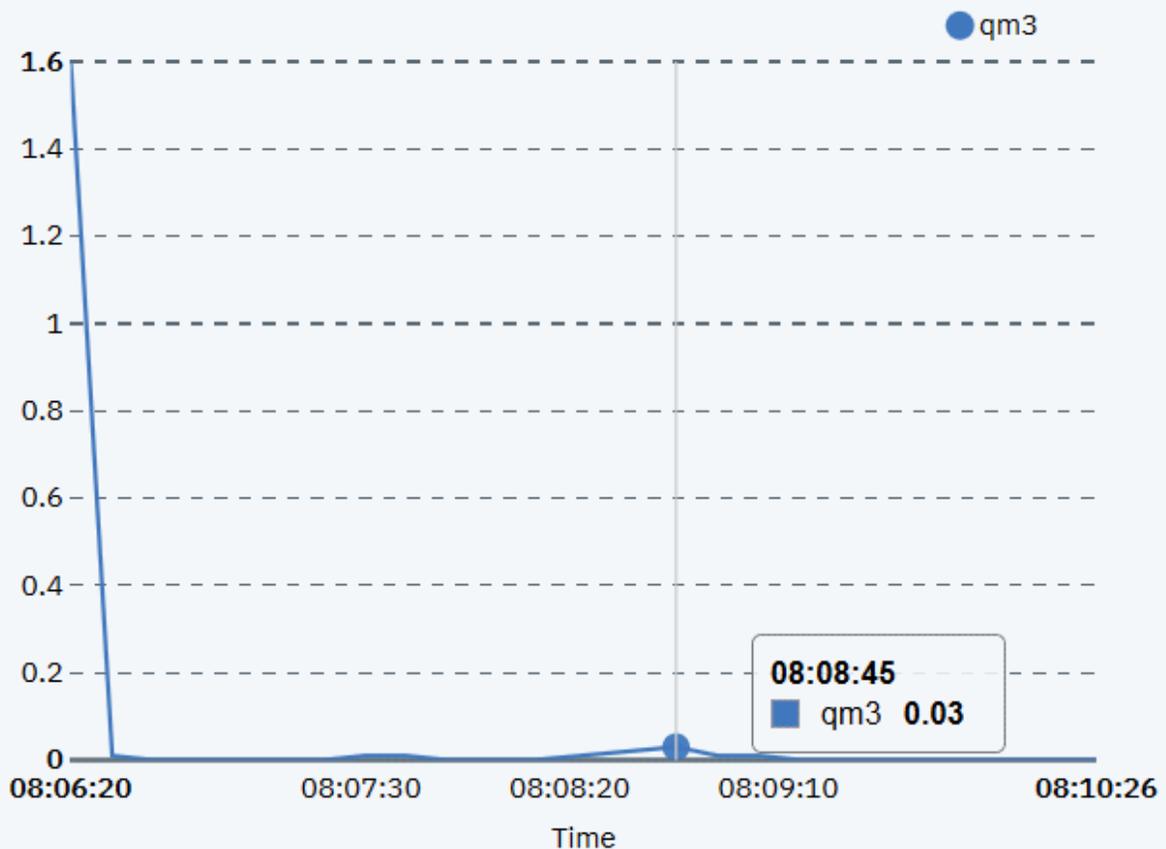
<i>Table 36. Elements for API per-queue usage statistics resources (continued)</i>		
Type	Element	Description
	MQGET browse non-persistent message count	Shows a count of non-persistent messages that have been browsed.
	MQGET browse persistent message count	Shows a count of persistent messages that have 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	

- f) 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.
- g) Click **Save**.

Results

After you configure the widget, there is a short delay before data is displayed in the chart. Data is displayed along a time axis. Each data point represents the end of the 10-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:

User CPU time - percentage estimate for que...



Interpreting CPU monitoring information

You can use the CPU monitoring information to see how well used 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% used. 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 used, 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%]
```

Configuring dashboard layouts

A dashboard is a container in the IBM MQ Console in which widgets are shown. You can create multiple dashboard tabs to show different selections of information.

About this task

You can configure each dashboard tab by clicking the arrow next to the tab name . You can change the tab name, and add a description for the tab. You can also configure how many columns the tab has.

You can configure the layout of the widgets within a dashboard tab by dragging and dropping the widgets.

Creating and deleting dashboard tabs

You can automatically create a dashboard tab that shows information about a specific local queue manager. You can manually create and delete dashboard tabs.

About this task

When you automatically create a dashboard tab to show information about a specific local queue manager, the following widgets are automatically added:

- Queues widget
- Client-connection channels widget
- Channels widget
- Listeners widget
- Subscriptions widget
- Topics widget
- Authentication information widget

Procedure

- To create a dashboard tab:

Tab 4 Tab 5 ▼ +

- Click the plus icon next to your existing dashboard tabs .
 - Enter a name for the new tab.
 - Optional: Enter a description for the new tab.
 - Click **Add**.
- To automatically create a dashboard tab for a specific queue manager:
 - Select the queue manager in the local queue manager widget.
 - Select **... > Add new dashboard tab**
A new dashboard tab is created. The tab has the name of the queue manager.
 - To delete a dashboard tab:

- Click the arrow next to the dashboard tab name .
- Select **Delete tab**.
- Confirm that you want to delete the dashboard tab by clicking **Delete**. The tab is deleted.

Importing and exporting dashboard layouts

You can save a dashboard layout by exporting it from the IBM MQ Console. You can import a saved dashboard layout into the IBM MQ Console.

About this task

When you export a dashboard, you create a .json file on your local disk. Subsequently, you can import the .json file to a dashboard to re-create the layout. When you import a dashboard layout, you can choose to add the imported tabs to an existing dashboard layout. Alternatively, you can replace the existing dashboard layout with the imported layout.

Procedure

- To export a dashboard layout:

- Click the dashboard menu icon .
- Select **Export dashboard**.
The file is saved to your browser download folder.

- To import a dashboard layout:

- Click the dashboard menu icon .
- Select **Import dashboard**.
The Import Dashboard Configuration window opens.
- Click **Browse** and browse for the location of the file that contains your configuration.
- Choose how to import the dashboard tabs:
You can choose from the following options:
 - **Append imported dashboard tabs to existing dashboard**
 - **Replace existing dashboard with imported dashboard tabs**
- Click **Import**.

The dashboard tabs are imported.

The dashboard controls

Use the controls at the top of the dashboard to enable IBM MQ Console tracing, access the online help, view information about the IBM MQ Console, and log out of the IBM MQ Console.

About this task

The dashboard controls apply to general features of the IBM MQ Console.

Procedure

- Access the IBM Knowledge Center for IBM MQ by clicking the help icon .
- Click the menu icon  to import, export, or reset the dashboard. For more information, see [“Configuring dashboard layouts” on page 344](#).
- Click the settings icon  to enable or disable diagnostic trace for the IBM MQ Console.
- View information about the IBM MQ Console, including information about the user that is logged in, by clicking the person icon  and selecting **About**.

Keyboard shortcuts

You can use keyboard shortcuts to work with the IBM MQ Console.

The following table lists the available shortcuts. All of the shortcuts are used from within a widget.

Key	Action
p	View properties of selected object
c	Create a new object
shift-d	Delete the selected object
a	Select all the objects in a widget
shift-a	Deselect objects

Administering IBM MQ by using the REST API

You can administer IBM MQ objects, such as queue managers and queues, by using the administrative REST API. Information is sent to, and received by, the administrative REST API in the JSON format.

The default URL to access the administrative REST API is:

`https://localhost:5554/ibmmq/rest/v1/admin`

The port number is the same as that used by the appliance REST management interface, see [“Configuring the REST management interface” on page 244](#).

To use the administrative REST API, you must be a user with access to the `mq/webadmin` or `mq/webuser` resource, and read access to the `login/rest-mgmt` resource (see [“User authorization, credential mapping, and access profiles” on page 423](#) and [“Access policies” on page 425](#)).

For detailed information, see [Using the administrative REST API](#) in the IBM MQ documentation.

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
“addmqm (add existing queue manager)” on page 644	Add an existing queue manager that uses SAN storage (used for recovery).
“crtmqm (create queue manager)” on page 578	Create a queue manager.
“dltmqm (delete queue manager)” on page 581	Delete a queue manager.
“dmpmqcfg (dump queue manager configuration)” on page 582	Dump the configuration of a queue manager.
“dspmq (display queue managers)” on page 585	Display information about queue managers.
“endmqm (end queue manager)” on page 597	Stop a queue manager.
“rmvmqinf (remove queue manager that uses SAN)” on page 658	Remove a queue manager that uses SAN storage. (SAN storage is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.)
“runmqsc (run MQSC commands)” on page 606	Run MQSC commands on a queue manager.
“setmqsize (increase queue manager file system size)” on page 659	Increase the size of the file system allocated to a queue manager.
“strmqm (start queue manager)” on page 610	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 25](#).

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 [“Types of user and how they are authenticated” on page 423](#) 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.

Messages that are generated by a queue manager, rather than an application, have the MQMD.UserIdentifier for that message set to mqsystem. Examples of situations where a message is generated by a queue manager are cluster internal state messages being sent to a remote queue manager and distributed pub/sub internal state messages being sent to a remote queue manager. In certain circumstances you might need to know that MQMD.UserIdentifier is set to mqsystem, for example, if messages generated by a queue manager are sent over a channel whose receiving side is using PUTAUT(CTX).

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 (create user)” on page 624	Creates user IDs for messaging users on the IBM MQ Appliance.
“userdelete (delete user)” on page 625	Deletes messaging users.
“usermodify (modify user)” on page 625	Modifies messaging users
“userlist (list users)” on page 626	Lists messaging users, or lists details of a particular user ID.
“groupcreate (create user group)” on page 626	Adds user groups for messaging users on the IBM MQ Appliance.
“groupdelete (delete user group)” on page 627	Deletes groups for messaging users.
“grouplist (list user groups)” on page 627	Lists groups for messaging users.
“userbackup (back up users)” on page 627	Backs up messaging users on the IBM MQ Appliance to a file.
“userrestore (restore user)” on page 628	Restores messaging users on the IBM MQ Appliance from a file to which they were previously backed up.

Related concepts

[“Messaging user and group commands” on page 624](#)

Appliance administrative users use these commands to set up user IDs and groups for messaging users.

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 `mqccli` prompt. See [“runmqsc \(run MQSC commands\)” on page 606](#) 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 31](#) 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 V9.1 Clients SupportPac from the following link: [MQC91: IBM MQ 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 clients, see [Installing and uninstalling IBM MQ on Linux and Installing and uninstalling IBM MQ on Windows](#) 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 \(run MQSC commands\)” on page 606](#).

Procedure

1. Type `mqccli` 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 347](#)). This user ID is the authority under which the client connection runs on the queue manager.

For example:

```
usercreate -u testuser -p passwd
```

3. Create and start a queue manager (see [“Message queue control commands” on page 347](#)):

```
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 V9 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 349.
- 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.

Related tasks

[Working with local queue managers](#)

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 349.
- 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 349.

Follow the related link for information on publishing and subscribing to topics from the IBM MQ Console.

Related tasks

[“Working with topics”](#) on page 317

You can use the topics widget in the IBM MQ Console to add and delete topics, and view and set the properties of a topic.

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 349.
- 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.

Related tasks

[“Working with topics”](#) on page 317

You can use the topics widget in the IBM MQ Console to add and delete topics, and view and set the properties of a topic.

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_name -t clchltab
```

A file with the name `queue_manager_name_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 402.

Example client connection using CCDT

An example of creating a client connection on the appliance and downloading the CCDT to the client.

A queue manager stores client connection channel information in a client channel definition table (CCDT). This table is updated whenever a client connection channel is defined or altered. An IBM MQ client application uses the client channel definition table (CCDT) to determine the channel definitions and

authentication information to use to connect to a queue manager. For a client application to use the CCDT, it must either be copied to the client machine or the file should be placed in a location from where the client can access it. On the appliance, the file needs to be copied to a directory from where it can be downloaded.

The following example shows you how to create a CCDT and copy it.

1. Run the following commands at the MQCLI prompt on the appliance to create the required IBM MQ objects and a user account under which the client connection is established:

```
usercreate -u mquser -p passw0rd
crtmqm QM1
strmqm QM1
runmqsc QM1
  DEFINE QLOCAL('Q1')
  DEFINE CHANNEL('CLNTCHL') CHLTYPE(SVRCONN) TRPTYPE(TCP)
  DEFINE CHANNEL('CLNTCHL') CHLTYPE(CLNTCONN) QMNAME('QM1') CONNAME('<IP_ADDRESS>(2323)')
  TRPTYPE(TCP)
  DEFINE LISTENER('L1') TRPTYPE(TCP) PORT(2323)
  START LISTENER('L1')
END
```

2. Set a channel authentication rule to accept connection from the IP address of your client system:

```
runmqsc QM1
  SET CHLAUTH ('CLNTCHL') TYPE(ADDRESSMAP) ADDRESS(<client_IP_ADDRESS>) MCAUSER('mquser')
```

3. Grant access to the user account 'mquser' so that it can connect to and inquire on the queue manager, and also has the authority to put, get, inq, and browse the queue Q1:

```
SET AUTHREC OBJTYPE(QMGR) PRINCIPAL('mquser') AUTHADD(CONNECT,INQ)
SET AUTHREC PROFILE('Q1') OBJTYPE(Queue) PRINCIPAL('mquser') AUTHADD(PUT,GET,INQ,BROWSE)
```

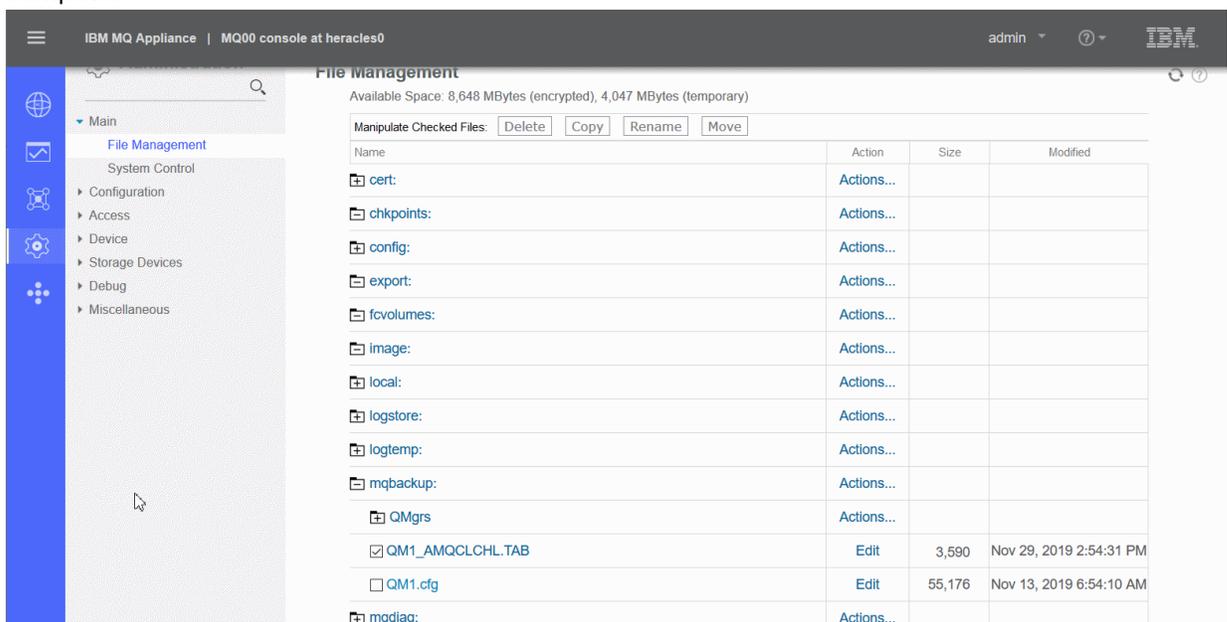
The next task is to create a file containing the CCDT in the appliance mqbackup: directory.

1. Run the following command on the appliance to generate the AMQCLCHL.TAB file in the mqbackup: // URI:

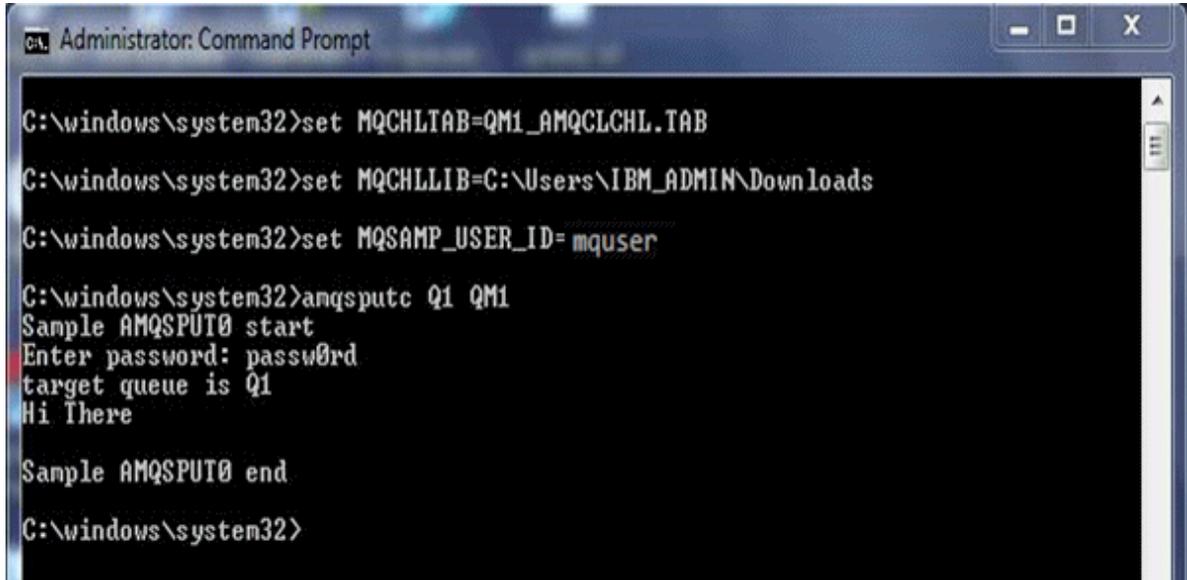
```
rcrmqobj -m QM1 -t clchltab
```

A file named QM1_AMQCLCHL.TAB is created in mqbackup: //.

2. Use the appliance web UI to download the CCDT file to the directory of your choice on the client computer.



3. You can take the following steps to run a sample application using the downloaded CCDT on the client computer:
 - a. On the Windows client computer, set the MQSAMP_USER_ID environment variable to identify the user that is running the sample programs.
 - b. Set the MQCHLLIB and MQCHLTAB environment variables to specify the location and name of the CCDT.



```
C:\windows\system32>set MQCHLTAB=QM1_AMQCLCHL.TAB
C:\windows\system32>set MQCHLLIB=C:\Users\IBM_ADMIN\Downloads
C:\windows\system32>set MQSAMP_USER_ID=mquser
C:\windows\system32>amqspuqc Q1 QM1
Sample AMQSPUT0 start
Enter password: password
target queue is Q1
Hi There
Sample AMQSPUT0 end
C:\windows\system32>
```

Starting and stopping the appliance

You can shut down and restart the appliance by using the command line interface or the IBM MQ Appliance web UI.

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 can 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 207. Log in as an administrative user.
2. Type the following command:

```
shutdown reboot
```

You can also shut down and restart by using the IBM MQ Appliance web UI. Complete the following steps:

1. Start the web UI as described in [“Configuring the IBM MQ Appliance web UI”](#) on page 209.
2. Click the administration icon  and select **Main > System Control**.
3. Set the Shutdown **Mode** to **Reboot system**.
4. Click **Shutdown**. The appliance restarts.

Related reference

[“shutdown”](#) on page 828

This command restarts or shuts down the appliance.

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 207. Log in as an administrative user.
2. Type the following command:

```
shutdown poweroff
```

You can also shut down the appliance by using the IBM MQ Appliance web UI. Complete the following steps:

1. Start the web UI as described in [“Configuring the IBM MQ Appliance web UI”](#) on page 209.



2. Click the administration icon and select **Main > System Control**.
3. Set the Shutdown **Mode** to **Power off system**.
4. Click **Shutdown**. The appliance restarts.

Related reference

[“shutdown”](#) on page 828

This command restarts or shuts down the appliance.

Restarting queue managers by using the command line

You can configure queue managers to auto-start.

Queue managers can be configured to automatically start when the appliance is restarted, or they can be configured to only start when requested by an administrator. Queue managers that are configured for high availability restart automatically by default.

Use the **crtmqm** command with the **-sa** option to create a queue manager with the auto-start feature enabled (see [“crtmqm \(create queue manager\)”](#) on page 578).

To check the auto-start capability of an existing queue manager, use the following command:

```
dspmqini -m QM_name -s InstanceData
```

Where *QM_name* is the name of the queue manager whose auto-start setting you want to check. The command displays the InstanceData stanza of the *qm.ini* file. This contains the Startup key, which has the following setting:

- Startup = Automatic, auto-start is enabled.
- Startup = Manual, auto-start is disabled.

If **dspmqini** does not report the Startup key in the InstanceData stanza, then a setting of Manual is implied. This means that the queue manager does not start automatically when the appliance is restarted.

A high availability (HA) queue manager has a value of Startup = HA, and you cannot modify this. If the queue manager is removed from HA control, then the value of Startup is set to Manual (and you can set it to Automatic if required).

To change the auto-start setting of an existing queue manager, you must edit the InstanceData stanza of the *qm.ini* file to change the setting of Startup to the required state. You use the **setmqini** command to change the setting. The following command enables auto-start:

```
setmqini -m QM_name -s InstanceData -k Startup -v Automatic
```

The following command disables auto-start:

```
setmqini -m QM_name -s InstanceData -k Startup -v Manual
```

Related reference

[“crtmqm \(create queue manager\)” on page 578](#)

Create a queue manager.

[“dspmqini \(display mq ini\)” on page 617](#)

Display attributes from the `qm.ini` or `mqat.ini` file of a specified queue manager. Display attributes from the `mq.s.ini` file.

[“setmqini \(set mq ini\)” on page 620](#)

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. Set or delete a diagnostic message definition in the `mq.s.ini` file.

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 and data.
5. IBM MQ Appliance web UI configurations.

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. Restore key repository.
4. Restore queue manager configurations and data.
5. Restore IBM MQ Appliance web UI configurations.

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, if required.

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 207](#).
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:

- Connect to the command line of the appliance as described in [“Command line access” on page 207](#).
- Log in to the appliance as an administrator.
- Type `config` to enter configuration mode.
- 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:

- Start the IBM MQ Appliance web UI, and click the menu icon  in the title bar.
- Select **Files** to open the File Management window.
- Open the `config` folder.
- 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).

Related tasks

[“Restoring the appliance configuration” on page 359](#)

You can restore the configuration of an IBM MQ Appliance to the same or to a different appliance.

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

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 88.
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 207.
- 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 menu icon  in the title bar.
 - b) Select **Files** to open the File Management window.
 - c) Click **Actions** for the config folder.
 - d) Select **Upload files** from the **Actions** menu.
 - e) Click **Browse**, and browse for the location of the file on your local system.
 - f) 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 207.
 4. Shut down and restart the appliance by entering the following command:

```
shutdown reboot
```

5. Re-create any appliance user IDs that were previously configured. See [“Configuring appliance user access”](#) on page 243.

Related tasks

[“Backing up or saving the appliance configuration”](#) on page 358

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.

Related reference

[“copy”](#) on page 788

This command copies a file to or from the appliance.

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 locally defined appliance users. When you restore a system you must re-create local appliance users and groups manually.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.
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 207.
- 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 menu icon  in the title bar.
- b) Select **Files** to open the File Management window.
- c) Open the mqbackup folder.
- d) 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).

Related concepts

[“Types of user and how they are authenticated”](#) on page 423

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.

Related tasks

[“Restoring messaging users”](#) on page 361

You can restore messaging user accounts that you previously backed up.

Related reference

[“copy”](#) on page 788

This command copies a file to or from the appliance.

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.

The following rules are applied when messaging users are restored:

- All existing messaging user and group configurations in the target appliance are replaced by the user and group configurations from the backup file, with the exception of system users and groups, and the mqm user and group (which have special handling).
 - The system users include root, mqm, mqsystem, hacluster, and sshd.
 - The system groups include root, mqm, haclient, sshd, and utmp.
 - The membership of the mqm group on the target appliance will be replaced with the group membership of the mqm from the backup file. (The ID of the mqm group is never changed.)
- Any messaging users or groups that exist on the target appliance but are not present in the backup file are deleted from the target appliance.
- The **userrestore** command prompts for confirmation (Y/N) after processing the backup file, but before restoring the configuration. The number of users that will be deleted and added by the command is displayed.

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 207.
- 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 menu icon  in the title bar.
 - b) Select **Files** to open the File Management window.
 - c) Click **Actions** for the mqbackup folder.
 - d) Select **Upload files** from the **Actions** menu.
 - e) Click **Browse**, and browse for the location of the backup file on your local system.
 - f) 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 207.
 3. Type **mqcli** to enter IBM MQ configuration mode.
 4. Type the following command to restore the messaging users:

```
userrestore -f user_backup_filename
```

You are told how many users and groups will be deleted, and how many created, and are asked to confirm the operation. You can specify the **-F** (force) parameter to skip the user confirmation.

Related concepts

[“Types of user and how they are authenticated”](#) on page 423

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.

Related tasks

[“Backing up messaging users” on page 360](#)

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.

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

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 207](#).
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 207](#).
- 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 menu icon in the title bar.
- b) Select **Files** to open the File Management window.
- c) Open the mqbackup folder.
- d) 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).

Related tasks

[“Restoring a key repository” on page 364](#)

You can restore a queue manager key repository that you previously backed up.

Related reference

[“keybackup” on page 637](#)

Back up the queue manager key repository to a file.

[“copy” on page 788](#)

This command copies a file to or from the appliance.

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 207](#).
- 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 menu icon  in the title bar.
 - b) Select **Files** to open the File Management window.
 - c) Click **Actions** for the mqbackup folder.
 - d) Select **Upload files** from the **Actions** menu.
 - e) Click **Browse**, and browse for the location of the backup file on your local system.
 - f) 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 207](#).
 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.

Related tasks

[“Backing up a key repository” on page 363](#)

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.

Related reference

[“keyrestore” on page 638](#)

Restore a key repository

[“copy” on page 788](#)

This command copies a file to or from the appliance.

[“listcert \(list certificate\)” on page 639](#)

List the certificates that are held in the keystore of a specific queue manager.

[“detailcert \(detail certificate\)” on page 635](#)

Show detailed information about a certificate for a specific queue manager.

Backing up a queue manager

You can use the command line to back up a queue manager to an archive file on the appliance.

About this task

You connect to the IBM MQ Appliance by using the command line, and save the queue manager to a file. The queue manager configuration is saved, together with log files and queue data.

Before you back up your first queue manager, you must create the target directory for backup files, and allocate storage for it in the appliance RAID volume.

A backup of a high availability (HA) queue manager does not contain any HA configuration data, so if you restore the queue manager from a backup file it is restored as a stand-alone queue manager. Similarly, disaster recovery (DR) configuration data is not preserved when you back up a DR queue manager.

V 9.1.4 The auto-config and auto-cluster information for a queue manager that is part of a uniform cluster is not backed up automatically. You should back up all files ending in `.ini` or `.mqsc` in the

mqbackup directory (see [“Configuring uniform clusters - appliance-specific considerations”](#) on page 261).

You can back up any type of queue manager while it is running, but this requires sufficient unallocated space on the internal disk to contain a temporary snapshot of the queue manager. This space is not required for a stand-alone queue manager if it is stopped before the backup is taken. HA and DR queue managers are always backed up from an internal snapshot, however, and so always require unallocated space on disk regardless of whether they are running or not.

If you are backing up so that you can use an archive file to migrate the queue manager, or if you want to be able to restore a queue manager to the state it was in at a particular time, then you should stop the queue manager before you back it up.

If the queue manager is running when you run the **mqbackup** command, a warning message is displayed.

If a queue manager is stopped before you take a backup, it is locked during the backup and cannot be started, deleted or otherwise changed while the backup runs.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.
2. Log in as a user in the administrators group.
3. Type **mqcli** to enter IBM MQ configuration mode.
4. If this is the first time you have backed up any queue manager, type the following command to allocate storage for your backup:

```
createbackupfs -s size
```

Where *size* is the size of the space that is allocated in GB. A directory that is named `mqbackup:///QMgrs` is created and allocated that storage.

5. Type the following command to back up a queue manager:

```
mqbackup -m QM_name
```

Where *QM_name* is the name of the queue manager that you want to back up. The backup can take some time to run, during which period you cannot use the CLI.

By default the archive file is named *QM_name.bak*, but you can add the `-o outfile` argument to the **mqbackup** command to specify a file name, if required:

```
mqbackup -m QM_name -o outfile
```

Related tasks

[“Restoring a queue manager”](#) on page 367

You can use the command line to restore a queue manager from an archive file that was created when you backed up the queue manager.

Related reference

[“createbackupfs:create appliance storage location for back up”](#) on page 666

Allocate space for back up archive files on the appliance.

[“deletebackupfs: clear previously allocated back up space”](#) on page 666

Free the space that was previously allocated for back up archive files on the appliance.

[“mqbackup: back up queue manager”](#) on page 667

Back up a queue manager.

[“mqrestore: restore queue manager from back up”](#) on page 668

Restore a queue manager from a back up archive.

Restoring a queue manager

You can use the command line to restore a queue manager from an archive file that was created when you backed up the queue manager.

Before you begin

Ensure that the archive file for the queue manager that you want to restore is located in the `mqbackup:///QMGRS` directory on the appliance.

About this task

You use the `mqrestore` command to restore a queue manager, including all its log files and data, from a previously taken backup. The command cannot run if there is already a queue manager with the same name on the appliance. The archive file must be located in the `backupfs` location, `mqbackup:///QMGRS`

You can only restore one queue manager at a time.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#).
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 restore a queue manager from its backup file:

```
mqrestore -f filename
```

Where *filename* is the backup file located in the `mqbackup:///QMGRS`. The restoration might take some time, during which the CLI is not available.

Related tasks

[“Backing up a queue manager” on page 365](#)

You can use the command line to back up a queue manager to an archive file on the appliance.

Related reference

[“mqbackup: back up queue manager” on page 667](#)

Back up a queue manager.

[“mqrestore: restore queue manager from back up” on page 668](#)

Restore a queue manager from a back up archive.

[“createbackupfs:create appliance storage location for back up ” on page 666](#)

Allocate space for back up archive files on the appliance.

[“deletebackupfs: clear previously allocated back up space” on page 666](#)

Free the space that was previously allocated for back up archive files on the appliance.

[“copy” on page 788](#)

This command copies a file to or from the appliance.

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:
 - a) Connect to the command line of the appliance as described in [“Command line access” on page 207](#).
 - 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 mqwebui:///com.ibm.mq.webui.persistence/username.json scp://yourusername@yourip[/]yourdirectorypath
```

For example:

```
copy mqwebui:///com.ibm.mq.webui.persistence/admin.json scp://midtownjj@server00d//safe/store/admin.json
copy mqwebui:///com.ibm.mq.webui.persistence/billg.json scp://midtownjj@server00d//safe/store/billg.json
```

- To copy the file by using the IBM MQ Appliance web UI:
 - a) Start the IBM MQ Appliance web UI, and click the menu icon  in the title bar.
 - b) Select **Files** to open the File Management window.
 - c) Open the `mqwebui` folder.
 - d) 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).

Related tasks

[“Restoring IBM MQ Appliance web UI configuration data” on page 368](#)

You can restore configuration data for each user who uses the IBM MQ Appliance web UI on the IBM MQ Appliance.

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

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:
 - a) Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#).
 - b) Log in as a user in the administrators group.
 - c) Type the following command to enter configuration mode:

```
config
```

d) 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.persistence
copy scp://midtownjj@server00d//safe/store/billg.json mqwebui:///
com.ibm.mq.webui.persistence
```

e) Type **exit** to leave config mode.

- To copy the file to the appliance by using the IBM MQ Appliance web UI:



- a) Start the IBM MQ Appliance web UI, and click the menu icon in the title bar.
- b) Select **Files** to open the File Management window.
- c) Click **Actions** for the mqwebui folder.
- d) Select **Upload files** from the **Actions** menu.
- e) Click **Browse**, and browse for the location of the user configuration file on your local system.
- f) 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.

Related tasks

[“Backing up IBM MQ Appliance web UI configuration data” on page 367](#)

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.

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

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 by using the command line” on page 174](#) 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 207](#). Log in as the user admin.
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](#).

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

[“reinitialize” on page 827](#)

This command deletes all configuration data from the file system of the appliance.

Triggering appliance operations by using the REST management interface

A subset of appliance operations can be triggered by using the REST management interface.

When you use the REST management interface for this purpose, you send HTTP requests to the REST interface port and receive JSON-formatted responses with a payload and indication of success or failure. You can incorporate requests into programs and so automate interaction with the appliance.

Before you trigger an operation on the appliance, you must become familiar with the `actionqueue` resource of the REST management interface. This resource exposes the appliance operation capabilities. You can find the available options for the `actionqueue` resource by sending a request like the one in the following example:

```
GET https://mqhost.com:5554/mgmt/actionqueue/
```

This request returns the following response, describing how to form an `actionqueue` request:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/actionqueue/"
    },
    "resource": {
      "href": "/mgmt/actionqueue/{domain}"
    },
    "operations": {
      "href": "/mgmt/actionqueue/{domain}/operations"
    },
    "schema": {
      "href": "/mgmt/actionqueue/{domain}/operations/{operation}"
    },
    "metadata": {
      "href": "/mgmt/metadata/{domain}/operations/{operation}"
    }
  }
}
```

To see the actual operations that are available for you to request, enter a request based on the following example:

```
GET https://mqhost.com:5554/mgmt/actionqueue/default/operations
```

The response that is returned from the `actionqueue` operations URI shows all supported operations on the appliance in the current firmware release. The following example shows a partial response payload to this request.

```
{
  "_links": {
    "self": {
      "href": "/mgmt/actionqueue/default/operations"
    },
    "AddPasswordMap": {
      "schema": {

```

```

    "href": "/mgmt/actionqueue/{domain}/operations/AddPasswordMap"
  },
  "metadata": {
    "href": "/mgmt/metadata/{domain}/operations/AddPasswordMap"
  }
},
"AddTrustedHost": {
  "schema": {
    "href": "/mgmt/actionqueue/{domain}/operations/AddTrustedHost"
  },
  "metadata": {
    "href": "/mgmt/metadata/{domain}/operations/AddTrustedHost"
  }
},
"ApplyPatch": {
  "schema": {
    "href": "/mgmt/actionqueue/{domain}/operations/ApplyPatch"
  },
  "metadata": {
    "href": "/mgmt/metadata/{domain}/operations/ApplyPatch"
  }
},
...
}

```

Search the response payload for the operation you want. If the operation does not appear in the returned list, the REST management interface does not currently support it. Instead, you must use one of the other interfaces to trigger that operation (for example, the CLI or the Web UI). If the operation is in the returned list, it is supported and can be triggered by using the REST management interface. You can use the schema and metadata embedded links to retrieve additional information about the required parameters for an operation. This is helpful when you compose the operation request payload, which is explained in the following section.

Compose a valid request payload

To trigger a supported operation by using the REST management interface, construct a valid request payload. You can use the metadata resource of the REST management interface to help you construct a valid payload. For example, to retrieve the metadata for the AddPasswordMap operation, you enter a request based on the following example:

```
GET https://mqhost.com:5554/mgmt/metadata/default/operations/AddPasswordMap
```

You receive the following response:

```

{
  "_links": {
    "self": {
      "href": "/mgmt/metadata/test-domain/operations/AddPasswordMap"
    },
    "doc": {
      "href": "/mgmt/docs/metadata/operations/AddPasswordMap"
    }
  },
  "action": {
    "name": "AddPasswordMap",
    "uri": "crypto/add-password-map",
    "cmd-group": "crypto",
    "cli-alias": "add password-map",
    "parameters": {
      "parameter": [
        {
          "name": "AliasName",
          "required": "true",
          "type": "dmString"
        },
        {
          "name": "Password",
          "required": "true",
          "type": "dmString"
        }
      ]
    },
    "display": "Add Password Map",
    "summary": "Add new password to password map",
    "description-encoded": "QWRkcyBhIHh3b3JkIHVvIHVvZSBlZGlzdGluZyBwYXN"
  }
}

```

```

        zd29yZCBtYXAuIFRoZSBwYXNzd29yZCBhbG1hcyBjYW4gYm
        UgdXNlZCBpbiB0aGUgY29uZm1ndXJhdGlvbiB0byByZWZlc
        iB0byB0aGUgcGFzc3dvcnQu"
    }
}

```

You can also acquire the metadata for the operation by looking up the appliance Service-Oriented Management Interface (SOMA) schema for the configuration object. The SOMA schemas are located in the store:///xml-mgmt.xsd file.

From the resource metadata you can identify the properties that are required by the operation. You can also identify the property names to use in the payload. By using this information, you can create a proper JSON request payload. A JSON payload has the following structure:

```

{
  "{operation_name}": {
    "{parameter1_name}": "{parameter1_value}",
    "{parameter2_name}": "{parameter2_value}",
    "{parameter3_name}": "{parameter3_value}",
    "{parameter4_name}": "{parameter4_value}"
    ...
  }
}

```

To produce a valid JSON request payload for the REST management interface, substitute `{operation_name}` and include all required parameters and their values. The following example shows a valid payload to trigger the `AddPasswordMap` operation.

```

{
  "AddPasswordMap": {
    "AliasName": "user",
    "Password": "passwd"
  }
}

```

Send a request and interpret the response

After the operation request payload is constructed, you can trigger the operation by using a POST request to the `actionqueue` resource. The following POST request shows how to trigger the `AddPasswordMap` operation, it is accompanied by the request payload:

```
POST https://mqhost.com:5554/mgmt/actionqueue/default
```

If the operation completed successfully, you see a confirmation response similar to the one in the following example:

```

{
  "AddPasswordMap": {
    "_links": {
      "self": {
        "href": "/mgmt/actionqueue/default"
      },
      "doc": {
        "href": "/mgmt/docs/actionqueue/operations/AddPasswordMap"
      }
    },
    "AddPasswordMap": "Operation completed."
  }
}

```

If the operation fails, an error is returned. To determine the cause of failure, examine the error message in the response payload and the appliance default log. One possible cause of failure is schema validation, as shown in the following example. This failure is caused by sending an improperly structured payload to trigger the operation.

```

{
  "_links": {
    "self": {
      "href": "/mgmt/actionqueue/default"
    }
  }
}

```

```
} ,
  "error": [
    "Schema validation failure. Please check the operation schema
    and default log for more information."
  ]
}
```

Related concepts

[“REST management interface” on page 1000](#)

The REST management interface provides access to the actions and to the configuration and status resources on 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 881](#).

Procedure

- To suspend an appliance in a high availability group, complete the following steps:
 - a) Log in to the appliance as a user in the administrators group.
 - b) Type the following command to enter IBM MQ administration mode:

```
mqcli
```

- c) Enter the following command:

```
sethagrp -s
```

- d) Enter the following command to check the appliance status:

```
dsphagrꞑ
```

The appliance is shown as being in the standby state.

- To resume the appliance after maintenance is finished, complete the following steps:
 - a) Log in to the appliance as a user in the administrators group.
 - b) Type the following command to enter IBM MQ administration mode:

```
mqcli
```

- c) Enter the following command:

```
sethagrꞑ -r
```

- d) Enter the following command to check the appliance status:

```
dsphagrꞑ
```

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.

Related reference

[“sethagrꞑ \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

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 881](#).

Procedure

1. Start the IBM MQ Appliance web UI and view the **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 stand-alone 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 272.
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 282.

Related reference

[“endmqm \(end queue manager\)”](#) on page 597

Stop a queue manager.

[“dlthagr \(delete HA group\)”](#) on page 672

Deletes a high availability (HA) group.

[“sethagr \(set HA group\)”](#) on page 675

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“strmqm \(start queue manager\)”](#) on page 610

Start a queue manager.

[“prepareha \(prepare appliance for HA group\)”](#) on page 674

Prepares an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.

[“crthagr \(create HA group\)”](#) on page 669

Creates a high availability (HA) group of appliances.

[“sethappreferred \(set HA preferred appliance\)”](#) on page 679

Sets a preferred appliance in the high availability (HA) group for a queue manager to run on.

[“clearhappreferred \(clear HA preferred appliance\)”](#) on page 679

Clears the preferred appliance for a queue manager in a high availability (HA) group.

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 **sethappreferred** 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 **sethappreferred** is used in conjunction with the **clearhappreferred** 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, **sethappreferred** 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 **sethappreferred** 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 **clearhappreferred** 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.

It is possible that conflicting **sethappreferred** and **clearhappreferred** commands might be issued when the appliances are disconnected. For example, both appliances might be set, or cleared, as the preferred location. When the appliances are reconnected, the conflict resolves to a consistent state, but that means that one of the commands on one of the appliances is silently undone. You should check which appliance is the preferred location when appliance reconnect after a fail over.

Procedure

- To set the current appliance as the preferred appliance for a queue manager:
 - a) Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```

- b) Specify that the current appliance is the preferred appliance for the named queue manager:

```
sethappreferred QMName
```

- To specify that the current appliance is no longer the preferred appliance for a queue manager:
 - a) Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```

- b) Specify that the current appliance is no longer the preferred appliance for the named queue manager:

```
clearhappreferred QMName
```

- To set the current appliance as the preferred appliance for a queue manager by using the IBM MQ Console:
 - a) Select the queue manager in the queue manager widget.
 - b) Select **More > High Availability** and select **Set preferred location**.
- To specify that the current appliance is no longer the preferred appliance for a queue manager by using the IBM MQ Console:
 - a) Select the queue manager in the queue manager widget.
 - b) Select **More > High Availability** and select **Clear preferred location**.

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 **dsphagr** command on the command line.

About this task

The **dsphagr** 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:

```
dsphagr
```

3. Optional: Exit the IBM MQ administration mode by entering the following command:

```
exit
```

Related reference

[“dsphagr \(display HA group status\)” on page 670](#)

Displays the status of the appliances in the high availability (HA) group.

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 view the **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 **sethagr -s** command).

Secondary appliance in standby mode

This status means that the other appliance in the HA pair has been suspended (by using the **sethagr -s** command).

Both appliances in standby mode

This status means that both appliances in the HA pair have been suspended (by using the **sethagr -s** command).

Secondary appliance unavailable

This status means that the connections to the other appliance in the HA pair have been lost.

Remote appliance(s) unavailable

This status means that the replication connection to the other appliance has been lost.

Partitioned

Queue manager data on the appliances is out of step, and cannot be automatically resolved.

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.

Inconsistent

The status is displayed on a secondary appliance during the initial synchronization of a queue manager if connection has been lost and synchronization was interrupted. The secondary appliance cannot provide high availability functionality until the initial synchronization has completed.

Resource failed

The physical interface associated with the floating IP used by the queue manager has been reconfigured. The queue manager is stopped and cannot be restarted until the original configuration is restored. This can also apply if an aggregated interface associated with the floating IP is changed.

This situation can also arise if the queue manager has been configured for disaster recovery and the eth20 interface has been reconfigured.

To recover, you must restore the original interface configuration, or remove the disaster recovery configuration. You can also delete the queue manager.

V 9.1.0 Migration Pending

When you have upgraded the appliance from Version 8.0 or Version 9.0 to Version 9.1, high availability queue managers display this status until you run the `sethagrp -u` command to upgrade them, and the other appliance is upgraded to Version 9.1 too. .

- 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 last synchronized. This information is shown when the status is `Partitioned` or `Inconsistent`.
- The time and date that the data was last in sync. This is displayed whenever data is out of sync.

Procedure

- To view the HA status of a queue manager by using the command line interface:
 - a) Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```
 - b) 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.
 - c) 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:
 - a) Open the console and find the widget that displays the queue manager.
 - b) Select the queue manager in the widget and select the properties icon from the toolbar  .
 - c) In the properties window, click on the **High availability status** section to open it.

Related reference

[“status” on page 881](#)

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.

Regenerating the keys for secure communication of the HA pair

Communication between the two appliance in a high availability (HA) pair is secured by private and public keys. You can check when the keys were last generated, and regenerate and exchange them if required.

About this task

The HA configuration sets up SSH communication between the two appliances on the primary interface (or, if that fails, the alternate interface). Public and private keys are created and exchanged when the HA pair are configured, but good security practice requires that these keys are periodically regenerated.

Procedure

- To regenerate keys by using the command line interface:
 - a) Enter the IBM MQ administration mode on either of the appliances by entering the following command:
`mqcli`
 - b) View date and time that the keys were last regenerated:
`dsphakeys`
The date and time are displayed in UTC time (UTC+00:00).
 - c) To regenerate the keys and exchange them, enter the following command:
`crthakeys`
(You cannot run this command if any other HA commands are running.)
 - d) Exit the IBM MQ administration mode by entering the following command:
`exit`
- To regenerate keys by using the IBM MQ Console:
 - a) Open the console and select **Manage HA keys** from the high availability menu on the title bar.
 - b) The date and time that the keys were last regenerated is displayed in UTC time (UTC+00:00) in the **Manage high availability SSH keys** window. Click **Regenerate keys** to regenerate and exchange the keys.

Related reference

[“crthakeys \(create HA keys\)” on page 670](#)

Regenerate SSH public and private keys and exchange public keys between appliances in an HA pair.

[“dsphakeys \(display HA key information\)” on page 671](#)

Display information about the SSH keys used for secure communication between an HA pair.

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 282](#).

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 381](#).

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 382](#).

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” on page 381](#) and [“Switching back to the main appliance by using the command line” on page 382](#). For steps to take following permanent failure, see [“Replacing a failed node in a disaster recovery configuration” on page 392](#).

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 399](#).

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 400](#).

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:
 - a) Log in to the recovery appliance as a user in the administrators group.

b) Type the following command to enter IBM MQ administration mode:

```
mqcli
```

c) 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.

d) Run the following command to start the queue manager:

```
strmqm QMname
```

- To start the secondary queue manager by using the IBM MQ Console:
 - a) Start the IBM MQ Appliance web UI and view the **MQ Console**.
 - b) In the queue manager widget, select the secondary queue manager that you want to make the primary.
 - c) Select **More > Disaster Recovery** and then select **Make DR primary**.
 - d) 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.

Related tasks

[Working with local queue managers](#)

Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“strmqm \(start queue manager\)” on page 610](#)

Start a 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

The first action that you must take in all circumstances is to restore the network connection between the main and recovery appliances and ensure that the connection is working.

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 561.

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, but synchronization is required between the main appliance and the recovery appliance.
- 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 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:
 - a) End the queue manager on the recovery appliance:

```
endmqm QMName
```

Where *QMName* identifies the queue manager.

- b) Specify that the queue manager on the recovery appliance is the secondary queue manager:

```
makedrsecondary -m QMName
```

- c) On the main appliance, specify that the original queue manager is now the primary once more:

```
makedrprimary -m QMName
```

- d) Ensure that 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.

- e) Start the queue manager on the main appliance:

```
strmqm QMName
```

This sequence is shown in diagrammatic form in Figure 1.

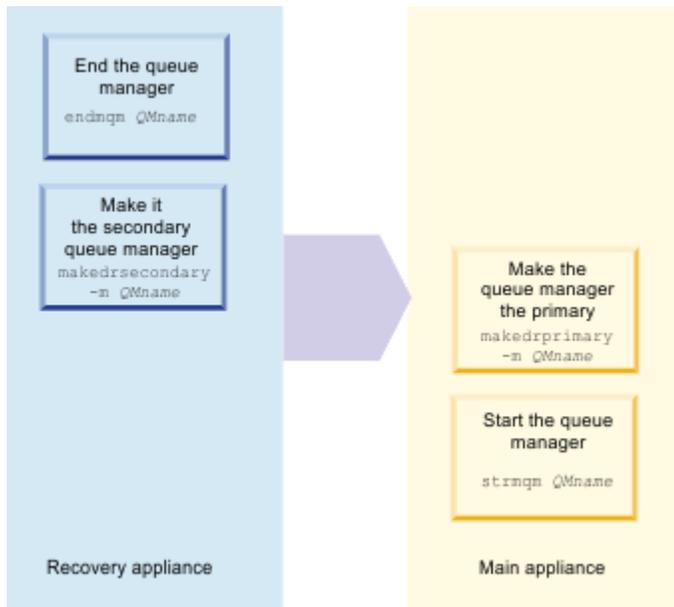


Figure 69. Switching back where there is no partitioning

- Where partitioning has occurred and you want to retain the data from the recovery appliance:
 - a) End the queue manager on the recovery appliance:

```
endmqm QMName
```

Where *QMName* identifies the queue manager.

- b) If required, end the queue manager on the main appliance:

```
endmqm QMName
```

- c) Specify that the queue manager on the main appliance is the secondary queue manager:

```
makedrsecondary -m QMName
```

- d) 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.

- e) 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.

- f) When the synchronization is complete, specify that the queue manager on the recovery appliance is now the secondary queue manager:

```
makedrsecondary -m QMName
```

- g) Specify that the queue manager on the main appliance is now the primary queue manager:

```
makedrprimary -m QMName
```

h) Start the queue manager on the main appliance:

```
strmqm QMName
```

This sequence is shown in diagrammatic form in Figure 2.

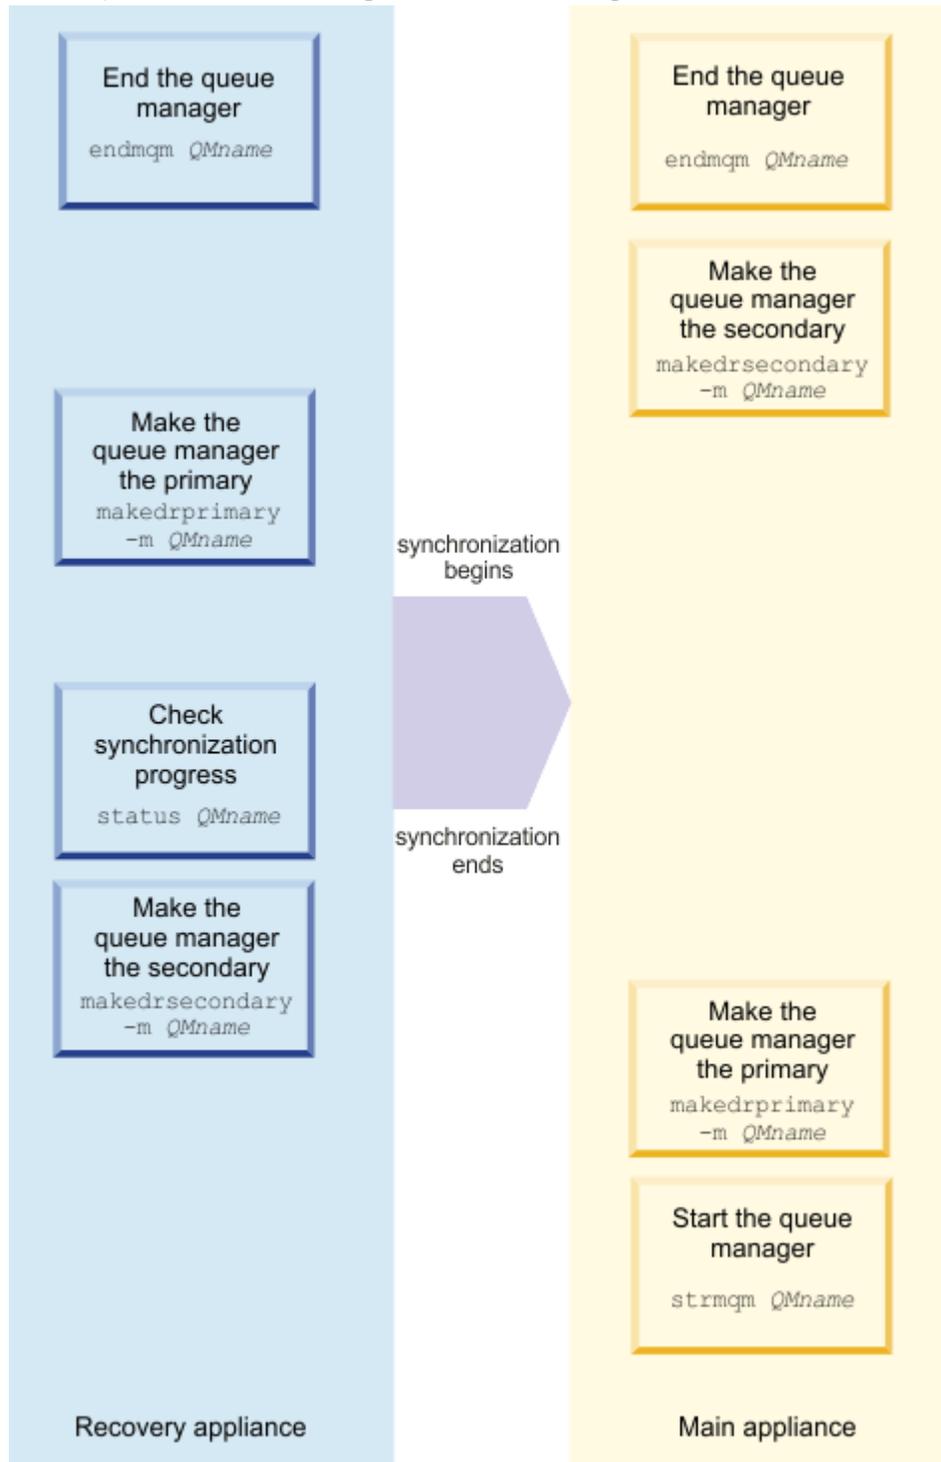


Figure 70. Switching back and retaining recovery appliance data

- Where partitioning has occurred and you want to retain the data from the main appliance:
 - a) End the queue manager on the recovery appliance:

```
endmqm QMName
```

Where *QMName* identifies the queue manager.

- b) If required, end the queue manager on the main appliance:

```
endmqm QMName
```

- c) Specify that the queue manager on the recovery appliance is now the secondary queue manager:

```
makedrsecondary -m QMName
```

- d) 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.

- e) 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.

- f) 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.

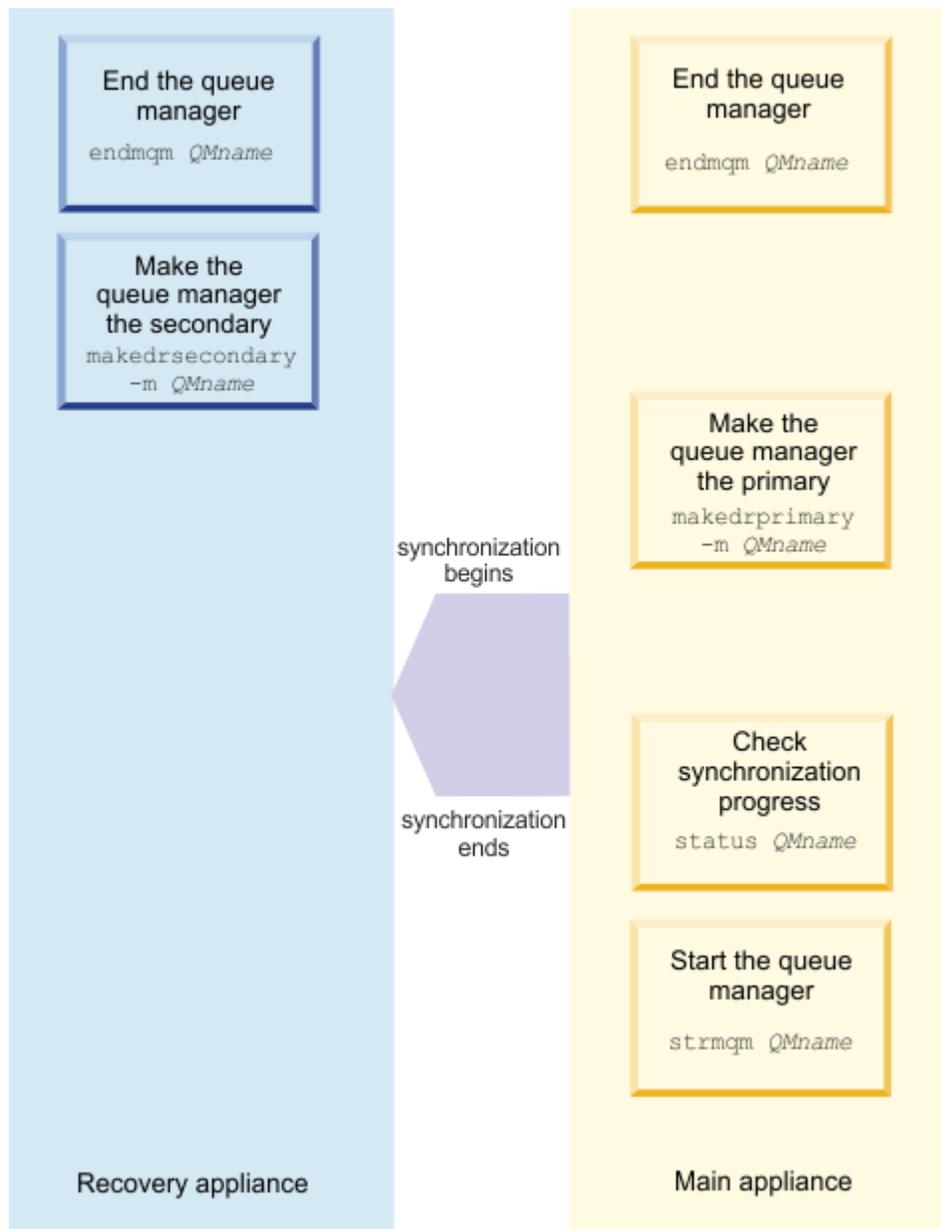


Figure 71. Switching back and retaining main appliance data

Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“makedrsecondary” on page 685](#)

Prevents a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.

[“status” on page 881](#)

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.

[“endmqm \(end queue manager\)” on page 597](#)

Stop a queue manager.

[“strmqm \(start queue manager\)” on page 610](#)

Start a queue manager.

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

The first action that you must take in all circumstances is to restore the network connection between the main and recovery appliances and ensure that the connection is working.

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 561.

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, but synchronization is required between the main appliance and the recovery appliance.
- 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 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:
 - a) Start the IBM MQ Appliance web UI on the recovery appliance and view the **MQ Console**.
 - b) In the queue manager widget, end the queue manager on the recovery appliance by selecting it and clicking the stop icon  .
 - c) Specify that the queue manager on the recovery appliance is the secondary queue manager. Select ... > **Disaster Recovery** and select **Make DR secondary**.
 - d) Ensure that synchronization has completed by selecting ... > **Disaster Recovery** and clicking **DR status**.
 - e) Start the IBM MQ Appliance web UI on the main appliance and view the **MQ Console**.
 - f) On the main appliance, specify that the original queue manager is now the primary once more by selecting ... > **Disaster Recovery (DR)** and clicking **Make DR primary**.
 - g) Start the queue manager on the main appliance by clicking the start icon  .

This sequence is shown in diagrammatic form in Figure 1.

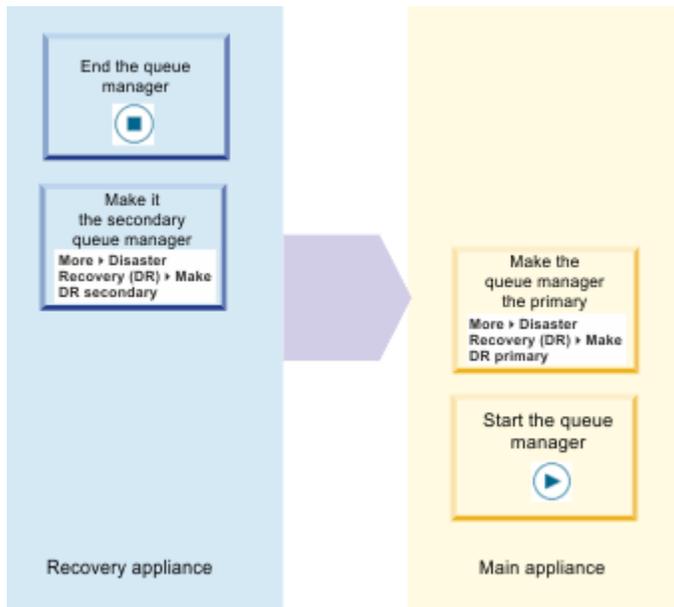


Figure 72. Switching back where there is no partitioning

- Where partitioning has occurred and you want to retain the data from the recovery appliance:
 - a) In the queue manager widget, end the queue manager on the recovery appliance by selecting it and clicking the stop icon .
 - b) If required, end the queue manager on the main appliance by selecting it and clicking the stop icon .
 - c) Specify that the queue manager on the main appliance is the secondary queue manager. Select ... > **Disaster Recovery**. and select **Make DR secondary**.
 - d) Specify that the queue manager on the recovery appliance is the primary queue manager. Select ... > **Disaster Recovery** and click **Make DR primary**. Synchronization begins.
 - e) Check whether the synchronization has completed by selecting ... > **Disaster Recovery** and clicking **DR status**.
 - f) When the synchronization is complete, specify that the queue manager on the recovery appliance is now the secondary queue manager. Select ... > **Disaster Recovery** and click **Make DR secondary**.
 - g) Specify that the queue manager on the main appliance is now the primary queue manager. Select ... > **Disaster Recovery** and click **Make DR primary**.
 - h) Start the queue manager on the main appliance by clicking the start icon .
- This sequence is shown in diagrammatic form in Figure 2.

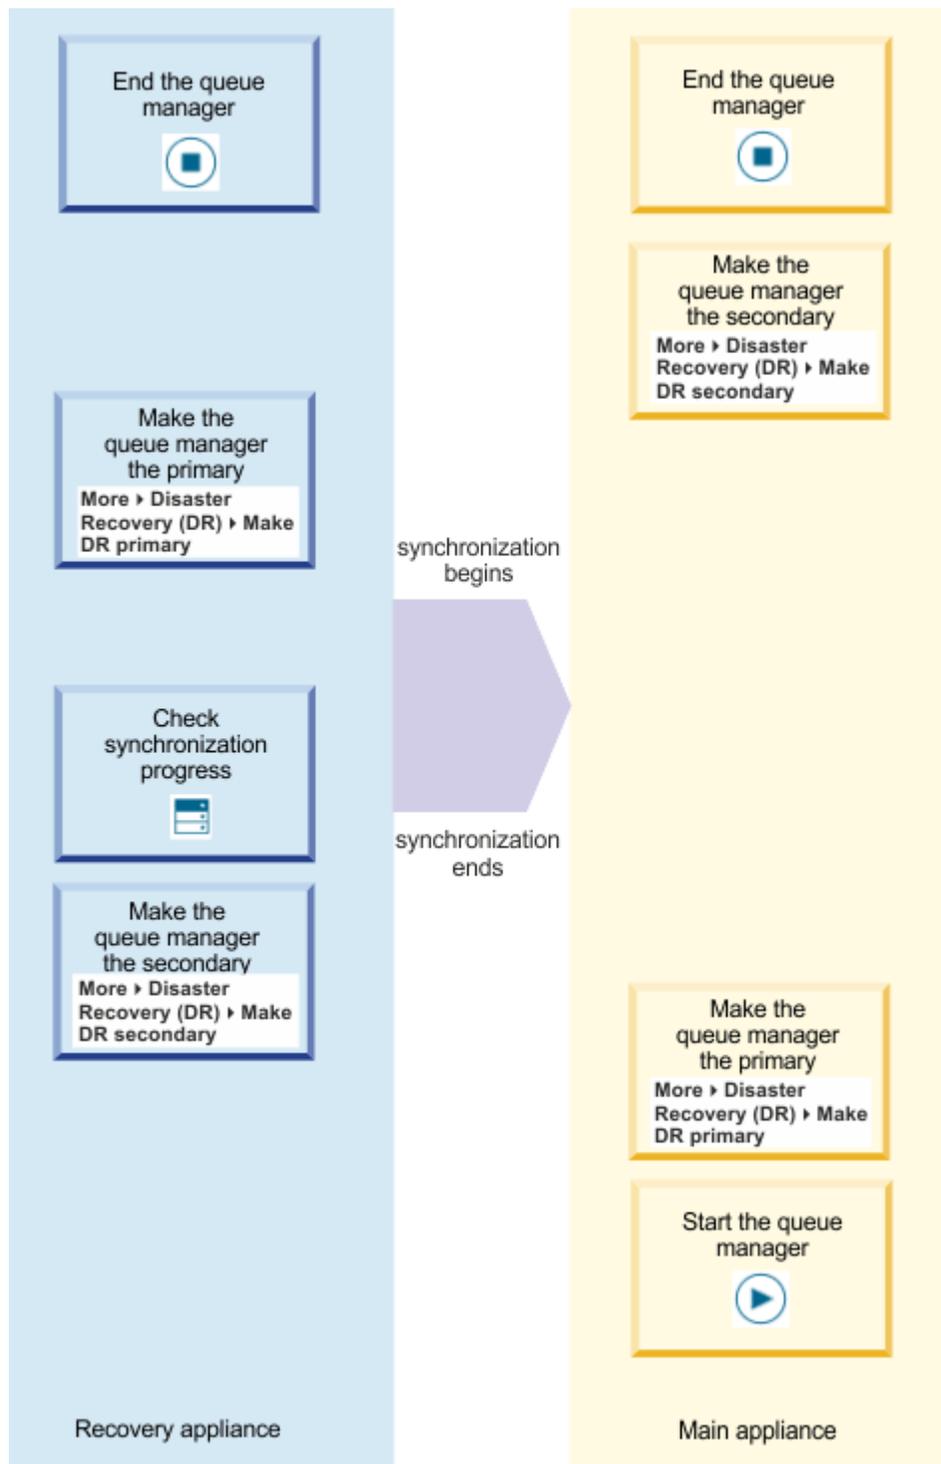


Figure 73. Switching back and retaining recovery appliance data

- Where partitioning has occurred and you want to retain the data from the main appliance:
 - a) In the queue manager widget, end the queue manager on the recovery appliance by selecting it and clicking the stop icon .
 - b) If required, end the queue manager on the main appliance by selecting it and clicking the stop icon .

- c) Specify that the queue manager on the recovery appliance is now the secondary queue manager. Select ... > **Disaster Recovery** and click **Make DR secondary**.
- d) Specify that the queue manager on the main appliance is now the primary queue manager. Select ... > **Disaster Recovery** and click **Make DR primary**. Synchronization begins.
- e) Check whether the synchronization has completed by selecting ... > **Disaster Recovery** and clicking **DR status**.
- f) 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): 

This sequence is shown in diagrammatic form in Figure 3.

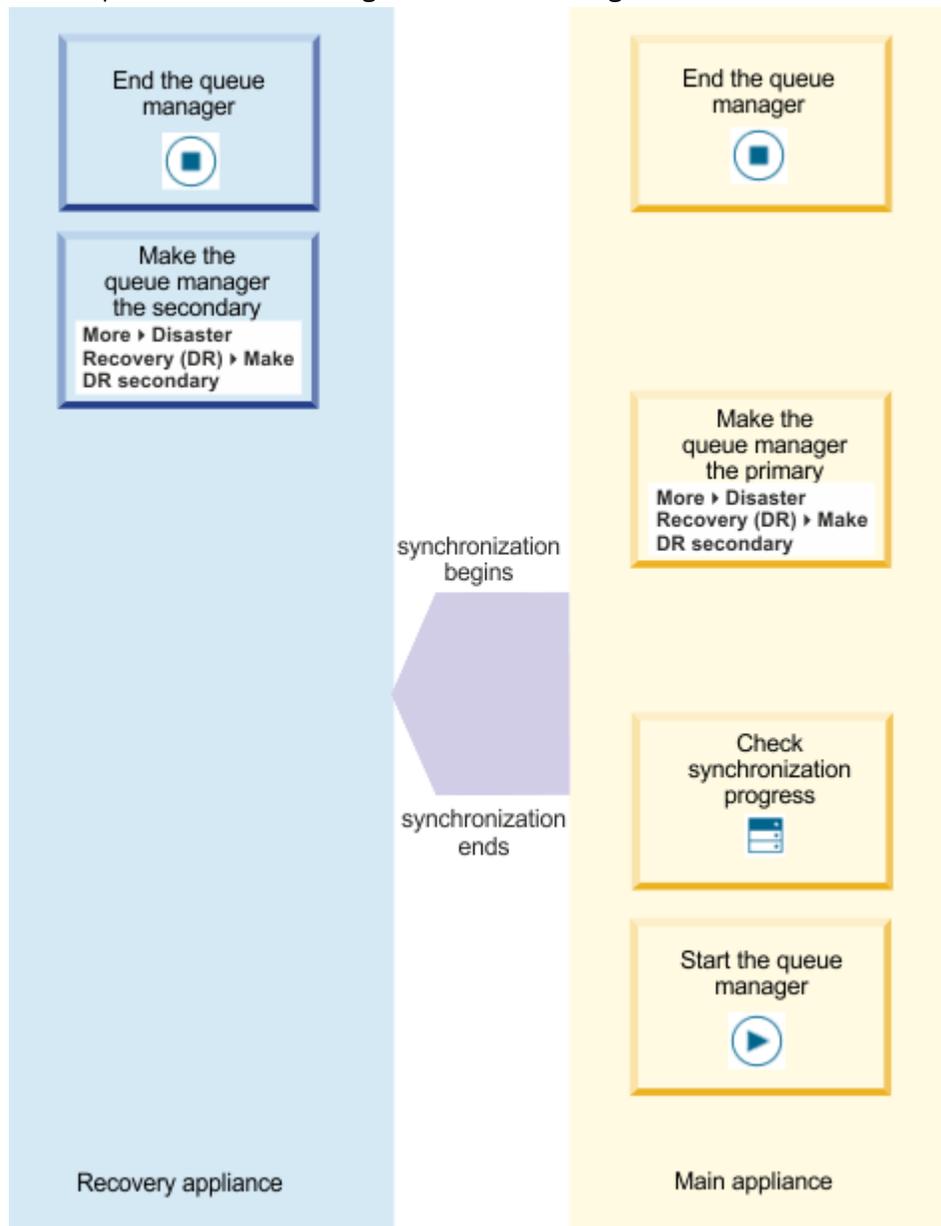


Figure 74. Switching back and retaining main appliance data

Related reference

“makedrprimary” on page 684

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

“makedrsecondary” on page 685

Prevents a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.

[“status” on page 881](#)

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.

[“endmqm \(end queue manager\)” on page 597](#)

Stop a queue manager.

[“strmqm \(start queue manager\)” on page 610](#)

Start a queue manager.

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:

```
dltldrprimary -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 288](#).
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:

```
makedrsecondary -m QMName
```

10. On the replacement appliance, make the queue manager the primary queue manager:

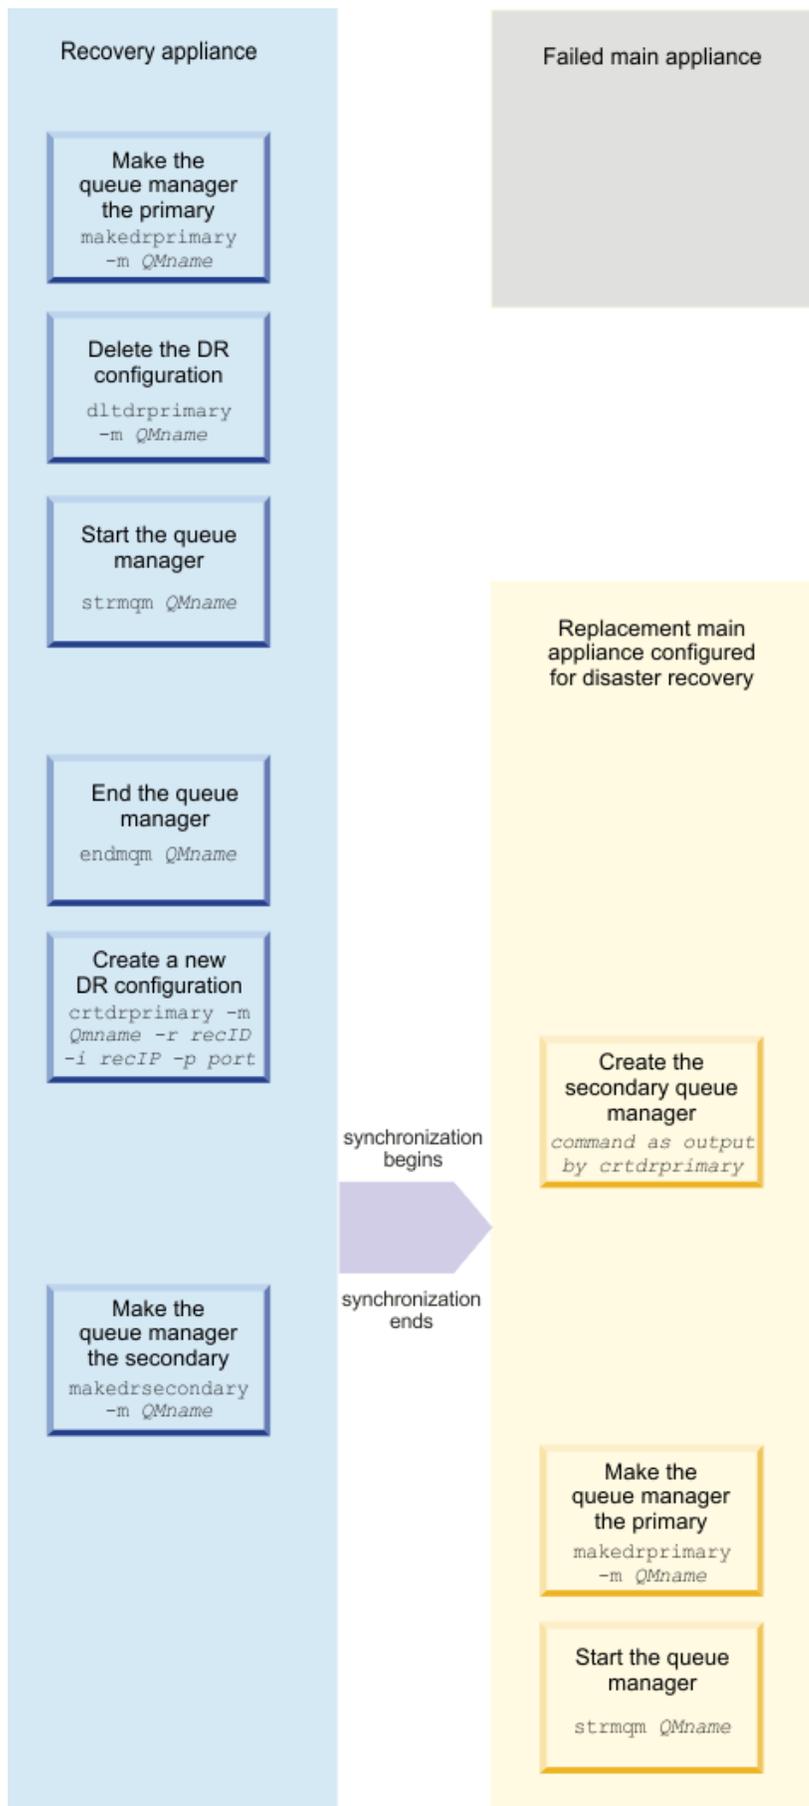
```
makedrprimary -m QMName
```

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:



Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“makedrsecondary” on page 685](#)

Prevents a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.

[“dltdrprimary” on page 682](#)

Remove a queue manager currently in the primary role from DR control.

[“crtldrprimary” on page 680](#)

Augments an existing queue manager to become the primary queue manager in a disaster recovery configuration.

[“status” on page 881](#)

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.

[“endmqm \(end queue manager\)” on page 597](#)

Stop a queue manager.

[“strmqm \(start queue manager\)” on page 610](#)

Start a queue manager.

Replacing failed high availability nodes in a disaster recovery configuration

If you lose high availability nodes which are also part of a disaster recovery (DR) configuration, you can replace the high availability appliances and restore the DR configuration by following this procedure.

About this task

If a disaster occurs such that the high availability appliances in the main site are beyond repair, you delete the disaster recovery configuration while your queue manager runs on the recovery appliance. You then replace the appliances, re-create the high availability group, and restore the disaster recovery configuration.

This procedure is shown in diagrammatic form in the following figure:

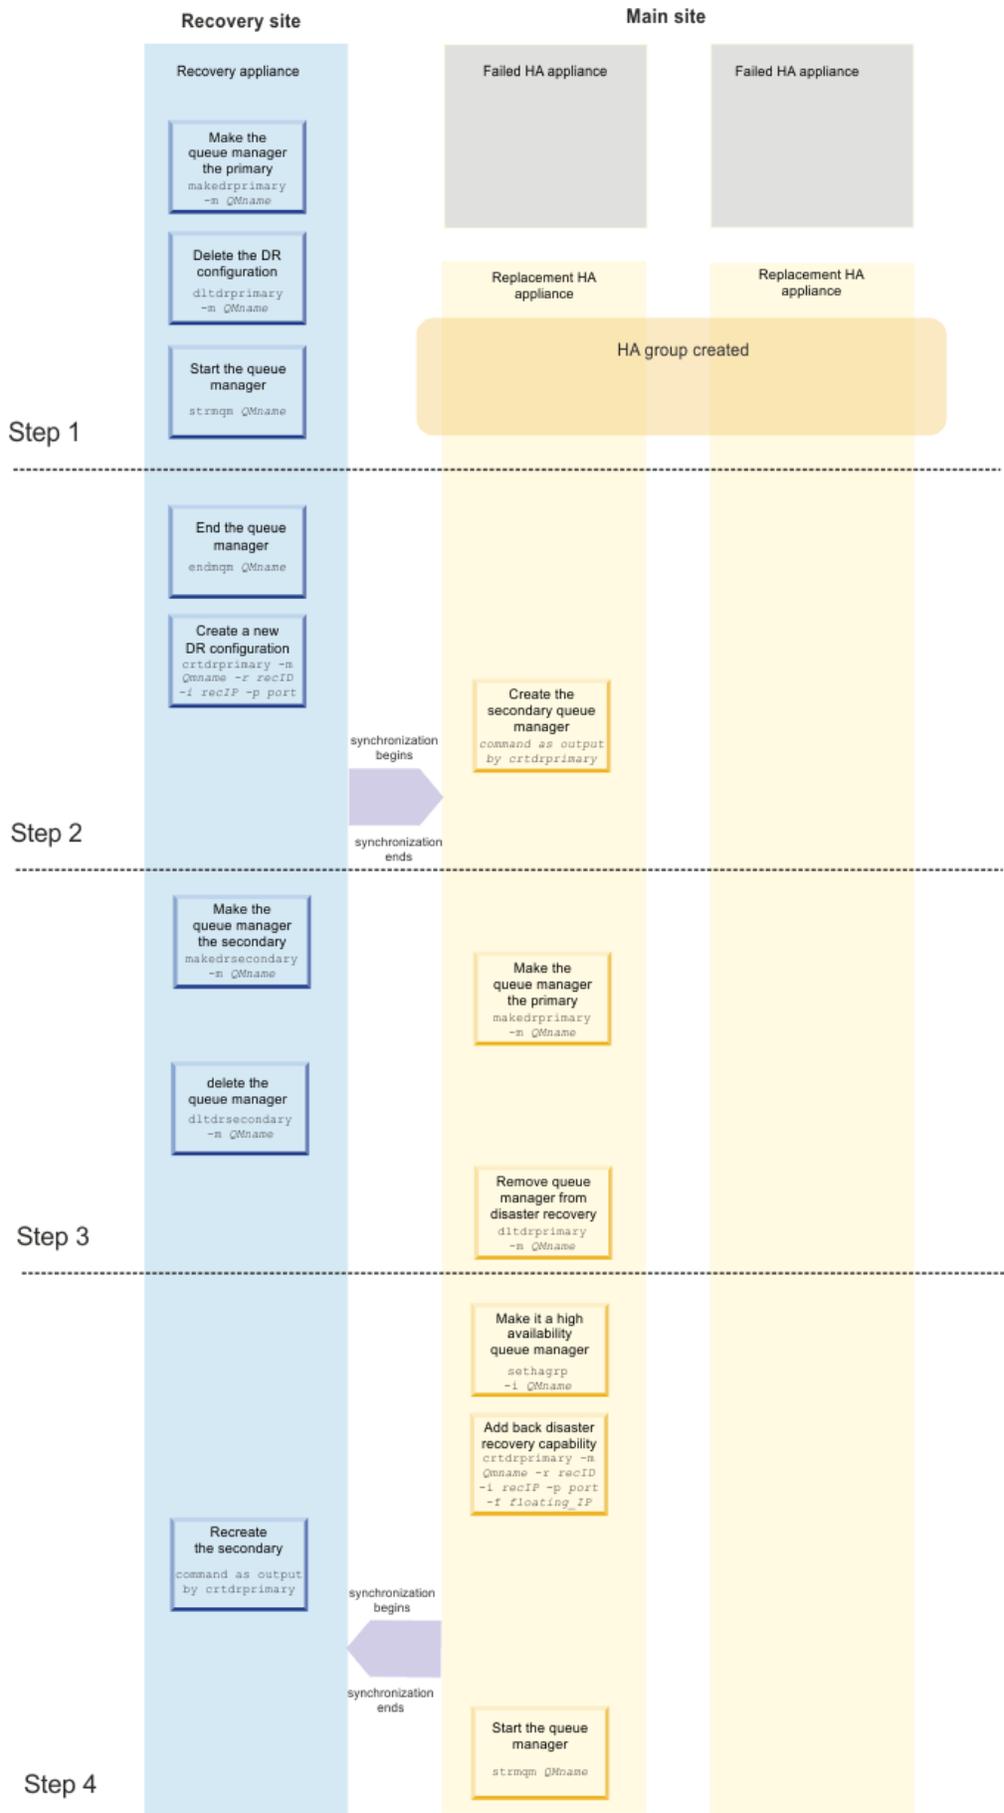


Figure 75. Procedure for replacing failed high availability nodes in a disaster recovery configuration

Procedure

Following the loss of the high availability queue manager on the main site, take the following steps:

1. Start running the queue manager on the DR recovery appliance and remove it from DR control to make it a stand-alone queue manager. At the main site, re-create the HA group on your replacement appliances:

- a) On the recovery appliance, run the following command:

```
makedrprimary -m QMname
```

Where *QMname* is the name of the queue manager.

- b) Delete the disaster recovery configuration:

```
dltdrprimary -m QMname
```

- c) Run the following command to start the queue manager as standalone:

```
strmqm QMname
```

- d) 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.
 - e) Replace the failed appliances on your main site and configure them for high availability, see [“Configuring the high availability group” on page 271](#).
2. Create a new DR configuration on your recovery appliance, with the queue manager running there as the primary. Use the output to create a secondary DR queue manager on your preferred main site appliance. At this point the queue manager resynchronizes and data is copied to the main site appliance.

- a) On the recovery appliance, stop the queue manager:

```
endmqm QMName
```

- b) On the recovery appliance, 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 that you want to be the preferred appliance for the queue manager.

-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

- c) On your preferred 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.

3. On the recovery appliance, make the queue manager the secondary and reestablish your main site appliance as the primary in the DR configuration. Then remove the queue manager from DR once again.

- a) On the recovery appliance, ensure that the queue manager is not running, then make it the secondary queue manager:

```
mkdrsecondary -m QMName
```

- b) On your preferred replacement appliance, make the queue manager the primary queue manager:

```
mkdrprimary -m QMName
```

- c) On the recovery appliance, delete the queue manager:

```
dltdrsecondary -m QMName
```

- d) On your preferred replacement appliance remove the queue manager from disaster recovery:

```
dltdrprimary -m QMName
```

4. Now re-create the high availability capability for the queue manager on the main site, and then add it back to the DR configuration.

- a) On your preferred replacement appliance, make the queue manager high availability, and then add it back to the disaster recovery configuration:

```
sethagip -i QMName  
crtldrprimary -m Qmname -r recID -i recIP -p port -f
```

The parameters for **crtldrprimary** are as described for step 7, with the addition of the following parameter:

-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.

- b) On the recovery appliance, re-create the secondary queue manager using the command that was output by the **crtldrprimary** command on the replacement appliance. Synchronization of data from the recovery appliance to the main appliance begins.
- c) On your preferred replacement appliance, start the queue manager:

```
strmqm QMname
```

Results

You have now restored the configuration as it was before the failure at your main site.

Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“makedrsecondary” on page 685](#)

Prevents a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.

[“dltdrprimary” on page 682](#)

Remove a queue manager currently in the primary role from DR control.

[“crtldrprimary” on page 680](#)

Augments an existing queue manager to become the primary queue manager in a disaster recovery configuration.

[“status” on page 881](#)

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.

[“endmqm \(end queue manager\)” on page 597](#)

Stop a queue manager.

[“strmqm \(start queue manager\)” on page 610](#)

Start a queue manager.

[“sethagr \(set HA group\)” on page 675](#)

► V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

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 the recovery appliance:

```
# ethernet eth20
Modify Ethernet Interface configuration
# admin-state disabled
# 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:

```
dldrprimary -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 enabled
# 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.

Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“dltdrprimary” on page 682](#)

Remove a queue manager currently in the primary role from DR control.

[“strmqm \(start queue manager\)” on page 610](#)

Start a queue manager.

[“endmqm \(end queue manager\)” on page 597](#)

Stop a queue manager.

[“dltmqm \(delete queue manager\)” on page 581](#)

Delete a queue manager.

[“admin-state” on page 708](#)

This command sets the administrative state for the configuration.

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.

Related reference

[“endmqm \(end queue manager\)” on page 597](#)

Stop a queue manager.

[“makedrsecondary” on page 685](#)

Prevents a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the 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 an in-progress synchronization has been interrupted. 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 last synchronized. This information is shown only when the status is `Partitioned`, `Inconsistent`, or `Remote appliance(s) unavailable`.
- The time and date that the data was last in sync. This is displayed whenever data is out of sync.
- 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:
 - a) Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```
 - b) 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.
 - c) 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:
 - a) Open the console and find the widget that displays the queue manager.
 - b) Select the queue manager in the widget. Select **More > Disaster Recovery** and click **DR status**.

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. Click the menu icon  and select **Files**.

File Management

Available Space: 10,728 MBytes (encrypted), 1,009 MBytes (temporary)

Manipulate Checked Files: Delete Copy Rename Move			
Name	Action	Size	Modified
cert:	Actions...		
chkpnts:	Actions...		
config:	Actions...		
export:	Actions...		
image:	Actions...		
local:	Actions...		
logstore:	Actions...		
logtemp:	Actions...		
mqbackup:	Actions...		
mqdiag:	Actions...		
mqerr:	Actions...		
mqpubcert:	Actions...		
mqmqdata:	Actions...		
mqtemporary:	Actions...		
mqtrace:	Actions...		
mqwebui:	Actions...		
policyframework:	Actions...		
pubcert:	Actions...		
sharedcert:	Actions...		
store:	Actions...		
tasktemplates:	Actions...		
temporary:	Actions...		

Manipulate Checked Files: Delete Copy Rename Move

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.

Managing files by using the REST management interface

You can use the REST management interface to manipulate files and directories on the IBM MQ Appliance.

When you use the REST management interface for this purpose, you send HTTP requests to the REST interface port and receive JSON-formatted responses with a payload and indication of success or failure. You can incorporate requests into programs and so automate interaction with the appliance.

File system navigation

To begin retrieving and modifying existing file system resources, you must become familiar with the filestore resource of the REST management interface that represents the appliance file system. You can find the format of the filestore resource by accessing the Uniform Resource Identifier (URI) of the filestore, as shown by the following example:

```
GET https://mqhost.com:5554/mgmt/filestore/
```

The following information is returned, which shows the required URI structure to manipulate individual files and directories on the appliance:

```
{
  "_links": {
    "self": {"href": "/mgmt/filestore/"},
    "top": {"href": "/mgmt/filestore/{domain}/{top_directory}"},
    "directory": {"href": "/mgmt/filestore/{domain}/{top_directory}/{directory_path}"},
    "file": {"href": "/mgmt/filestore/{domain}/{top_directory}/{file_path}"}
  }
}
```

Directory management

You can perform all directory manipulation operations. These operations include retrieving the contents of existing directories, creating directories and sub-directories, and deleting existing directories. As with all other REST requests on the IBM MQ Appliance, these requests specify the default domain.

Retrieving the contents of a directory

You can retrieve the contents of any appliance directory if you have appropriate access permissions to that directory. To retrieve the contents of any directory, construct a URI according to the `directory` link in the `filestore` resource listing. The following request shows an example in which the `local:///test-directory` directory is accessed in the default domain:

```
GET https://mqhost.com:5554/mgmt/filestore/default/local/test-directory
```

The response shows that the target directory contains one file, test-file, and the relevant information for that file.

```
{
  "_links": {
    "self": {
      "href": "/mgmt/filestore/default/local/test-directory"
    },
    "doc": {
      "href": "/mgmt/docs/filestore"
    }
  },
  "filestore": {
    "location": {
      "name": "local:/test-directory",
      "file": {
        "name": "test-file",
        "size": 1182,
        "modified": "2016-04-07 15:14:17",
        "href": "/mgmt/filestore/default/local/test-directory/test-file"
      },
      "href": "/mgmt/filestore/default/local/test-directory"
    }
  }
}
```

Create directories

You can create a directory with a PUT request or a POST request. Both requests accomplish the same operation, but require a different URI to complete successfully. You can choose which approach is more convenient for you. The following POST request shows the URI that is required to create a subdirectory in the `local:///` directory:

```
POST https://mqhost.com:5554/mgmt/filestore/default/local
```

The following PUT request shows the URI that is required to create a test-directory subdirectory within the `local:///` directory:

```
PUT https://mqhost.com:5554/mgmt/filestore/default/local/test-directory
```

Both the POST and PUT requests require that the details of the directory to be created are specified in the request payload. The following example shows the required request payload, with the directory name specified in the **name** parameter. This payload structure is used for both the PUT request and the POST request. The directory name in the payload and in the URI for the PUT request must match, otherwise an error results.

```
{
  "directory": {
    "name": "test-directory"
  }
}
```

Issuing a POST request or a PUT request on an existing directory resource returns an error. This feature protects you from accidentally removing the directory contents. If you intend to overwrite a directory with new contents, you must first delete the directory by issuing a DELETE request and then re-create it.

Delete existing directories

To delete an existing directory, send a DELETE request to the target directory. The following example request requests deletion of the `local:///test_dir` directory:

```
DELETE https://mqhost.com:5554/mgmt/filestore/default/local/test_dir
```

After the directory is deleted, you see a response similar to the following example:

```
{
  "_links": {
```

```

    "self": {
      "href": "/mgmt/filestore/default/local/test_dir"
    },
    "doc": {
      "href": "/mgmt/docs/filestore"
    }
  },
  "result": "ok",
  "script-log": ""
}

```

File management

You can perform all file manipulation operations by using the REST management interface. These operations include retrieving and updating the contents of existing files, creating files, and deleting existing files.

Retrieve file contents

You can retrieve the contents of any file on the appliance provided you have appropriate access permissions to that file. To retrieve the contents of any file, construct a URI based on the `file` part of the filestore resource. For example, the following request retrieves contents of the `test_file.txt` file in the `local:///test_dir` directory:

```
GET https://mqhost.com:5554/mgmt/filestore/default/local/test_dir/test_file.txt
```

File contents that are returned as a base64-encoded payload:

```

{
  "_links": {
    "self": {
      "href": "/mgmt/filestore/default/local/test_dir/test_file.txt"
    },
    "doc": {
      "href": "/mgmt/docs/filestore"
    }
  },
  "file": {
    "name": "local:///test_dir/test_file.txt",
    "value": "SEVMTE8hISE=..."
  }
}

```

Create and update files

You can create a file by using a PUT request or a POST request. Both requests create a file, but require a different URI to complete successfully. A POST request fails if a file with the same name exists in the target directory. This feature prevents you from accidentally overwriting an existing file. However, you can also create a file by using a PUT request. Issuing a PUT request on an existing file overwrites the file with the contents in the request payload.

The following POST request shows the URI that is required to create a file in the `local:///` directory:

```
POST https://mqhost.com:5554/mgmt/filestore/default/local
```

The following PUT request shows the URI that is required to create the `test-file.txt` file in the `local:///` directory:

```
PUT https://mqhost.com:5554/mgmt/filestore/default/local/test-file.txt
```

Both the POST and PUT requests require that the details of the file to be created are specified in the request payload. The following example shows the required request payload, with the file name specified in the `name` parameter and the contents in the `contents` parameter. The file contents must be base64-encoded before they are embedded into the request payload. This payload structure is

used for both the PUT request and the POST request. The file name in the payload and in the URI for the PUT request must match, otherwise an error results.

```
{
  "file": {
    "name": "test-file",
    "content": "dG9wOyBjb25maWd1cmUgdGVybWluYWw7CgojIGNvbmZpZ3VyYXRpb
      24gZ2VuZX4gOStAcmJowMludGVyZmFjZSAiZXRoMTAgYXJwCiAgagXB
      2NgogIG5vIHNSYWFjCiAgbXR1ICIxNTAwIgotIGlwdjYtbGlua2xv
      Y2FsLXN0YXJ0dXAtd28tcHJpbWFyeS1pcGFkZHIKICBhZG1pbi1zd
      GF0ZSAiZW5hYmxlZCICKXhpdCAKCludGVyZmFjZSAiZXRoMjAi..."
  }
}
```

After the file is created, a response is returned similar to the one shown in the following example:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/filestore/default/local/test-file"
    },
    "doc": {
      "href": "/mgmt/docs/filestore"
    }
  },
  "test-file": "File has been created."
}
```

If you use a PUT request to overwrite an existing file, you receive a response similar to the following example:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/filestore/default/local/test-file"
    },
    "doc": {
      "href": "/mgmt/docs/filestore"
    }
  },
  "test-file": "File has been updated."
}
```

Delete existing files

To delete an existing file, send a DELETE request to the target file. The following example shows this request for the `test_file.txt` file in the `local:///test_dir` directory:

```
DELETE https://mqhost.com:5554/mgmt/filestore/default/local/test_dir/test_file.txt
```

After the file is deleted, a response is returned that is similar to the following example:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/filestore/default/local/default/test_file.txt"
    },
    "doc": {
      "href": "/mgmt/docs/filestore"
    }
  },
  "result": "ok"
}
```

Related concepts

[“REST management interface” on page 1000](#)

The REST management interface provides access to the actions and to the configuration and status resources on the appliance.

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 831.
- You must create a special IPMI user, [“IPMI user commands”](#) on page 829.
- 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
```

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 V9.1 is compatible with IBM MQ V9.1. Follow these instructions if you are moving queue managers from an IBM MQ V9.1 platform.

You use the **dmpmqc.fg** command on your source system to save the configuration of a queue manager. Running **dmpmqc.fg** 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.

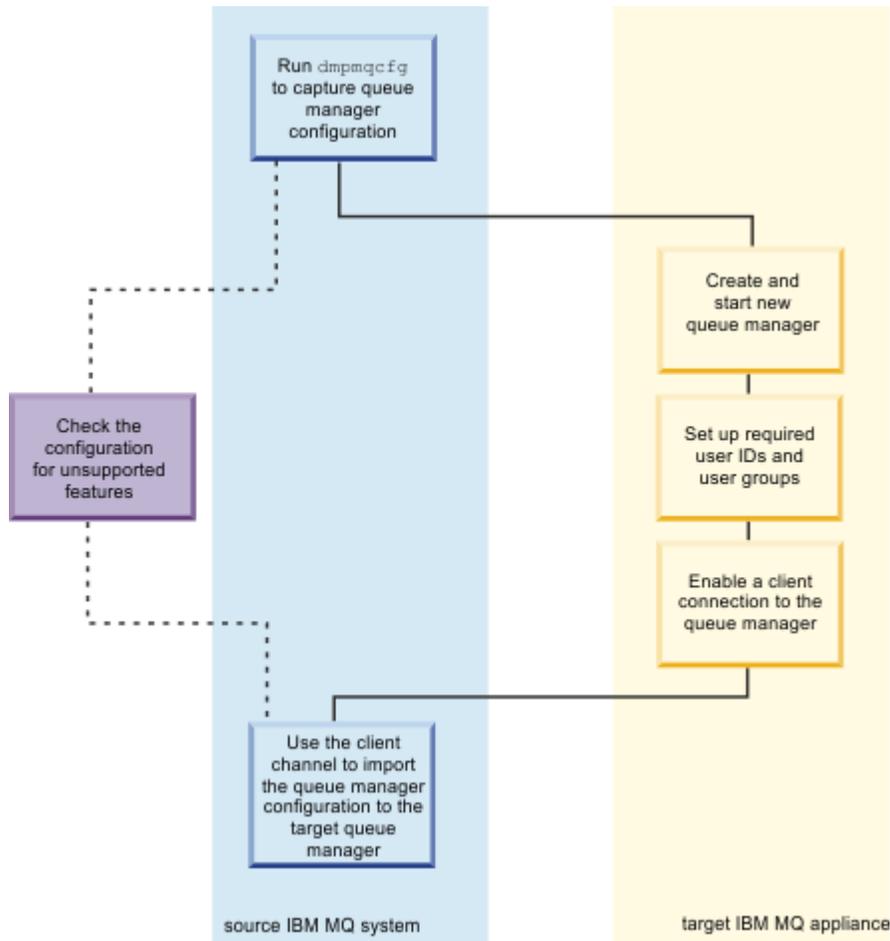
Moving a queue manager

Follow these instructions to move a queue manager from an IBM MQ V9.1 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:

```
dmpmqcfig -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.

Examine the `dmpmqcfig` file and mask out incompatible features as described in [“Planning for incompatible features in the queue manager”](#) on page 415 and [“Handling incompatible features in the queue manager”](#) on page 416.

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 311, or you can use MQ commands, with the required name and attribute values. If you want to use MQ commands, you must complete the following steps:
 - a) Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.
 - 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. For details about creating messaging users on the appliance, see [“Administering messaging users” on page 347](#).
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 311](#), 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 416](#). 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 on the source system to connect to the target system over the channel and replay on the target queue manager the commands that were exported from the source queue manager:

```
runmqsc -c target_QM_name < QM_file
```

9. Restore the attributes that were masked in the **dmpmqc fg** output and that you identified when you checked the output (as described in [“Handling incompatible features in the queue manager” on page 416](#)). 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
```

For example, if your **dmpmqc fg** output for your source queue manager contained the following entry:

```
DEFINE AUTHINFO('arthurauth') +  
  AUTHTYPE(CRLLDAP) +  
  CONNAME('william.paddock.com') +  
  LDAPUSER('joey') +  
 * LDAPPWD('*****') +  
 * ALTDATE(2019-11-13) +  
 * ALTTIME(06.41.13) +  
 REPLACE
```

You would need to specify the LDAP password that had been masked:

```
ALTER AUTHINFO('arthurauth')  
  LDAPPWD('ldap_serv_passw0rd')
```

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 207.
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=organization,C=country" -size size -expire expiry
```

For example:

```
mqa(mqcli)# createcert -m REGA -label ibmwebspheremqrega -dn "CN=Issuer,OU=Certificate  
Authority,O=IBM,C=GB" -size 1024 -expire 30
```

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_certificates/ibmwebspheremqregb.p12
```

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_certificates/  
ibmwebspheremqregb.p12 mqpubcert:///ibmwebspheremqregb.p12
```

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.

Related reference

[“copy” on page 788](#)

This command copies a file to or 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 25](#)
- [“Moving a queue manager” on page 411](#)
- [“Using an IBM MQ client ” on page 349](#)

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 **dmpmqc fg** output.

Special considerations for moving a queue manager from z/OS

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 **dmpmqc fg** 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 411 and [“Using an IBM MQ client”](#) on page 349.

Channels that use SSL 3.0 or TLS 1.0 CipherSpecs

By default, IBM MQ v9.1 does not support SSL 3.0 or TLS 1.0 and related CipherSpecs. If you move a queue manager to the IBM MQ Appliance that has one or more channels that use SSL 3.0 or TLS 1.0, you can take action to enable support for these CipherSpecs. You can take one of the following actions:

- Set the environment variable `AMQ_SSL_WEAK_CIPHER_ENABLE=ALL`. See [“Configuring environment variables”](#) on page 261 for details of how to set environment variables on the appliance.
- Change the SSL stanza in the `qm.ini` file to re-enable all SSL 3.0 and TLS 1.0 CipherSpecs, for example:

```
mqa# mqcli
mqa(mqcli)# setmqini -m QMGR -s SSL -k AllowWeakCipherSpec -v ALL
```

See [“Adding a value to the configuration file”](#) on page 417 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 `dmpmqcfig` 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 `dmpmqcfig` output, and amend it if necessary, to deal with incompatible features.

The output from the `dmpmqcfig` 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 `dmpmqcfig` 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 `dmpmqcfig`.

Substitute appropriate values for masked values

The output from the `dmpmqcfig` 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 **dmpmqcfig** 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 **dmpmqcfig** 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
- DiagnosticMessages
- InstanceData
- TuningParameters
- SSL
- Security
- Subpool
-  AutoConfig
-  AutoCluster
-  Variables

The following value for *Stanza* modifies the `mqat.ini` file.

- AllActivityTrace

The following values for *stanza* modify the `mqs.ini` file

- DiagnosticSystemMessages
- DiagnosticMessagesTemplate

For more information about Diagnostic Messages, see [“IBM MQ diagnostic message services” on page 505](#).

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 32](#).

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 RemoteQueueManagerAccessControl -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
```

Related reference

[“setmqini \(set mq ini\)” on page 620](#)

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. Set or delete a diagnostic message definition in the `mqc.ini` file.

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:

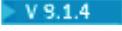
QMGrName

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
- DiagnosticMessages
- InstanceData
- TuningParameters
- SSL
- Security
- Subpool
-  AutoConfig
-  AutoCluster
-  Variables

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
```

Related reference

[“setmqini \(set mq ini\)” on page 620](#)

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. Set or delete a diagnostic message definition in the `mqsc.ini` file.

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.inifile`, 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
- DiagnosticMessages
- InstanceData
- TuningParameters
- SSL
- Security
- Subpool
-  AutoConfig

- [V 9.1.4](#) AutoCluster
- [V 9.1.4](#) Variables

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:

```
dspmqrini -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:

```
dspmqrini -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:

```
dspmqrini -m QM1 -s AllActivityTrace -k ActivityCount
```

Related reference

[“dspmqrini \(display mq ini\)” on page 617](#)

Display attributes from the `qm.ini` or `mqat.ini` file of a specified queue manager. Display attributes from the `mqs.ini` file.

Security

You must consider security requirements before and after you configure the IBM MQ Appliance.

Types of user and how they are authenticated

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

Authentication of appliance users, and authorization of them to access appliance resources, is controlled by role based management (RBM). RBM defines how users are authenticated and authorized. You can specify the following authentication methods:

- Users can be authenticated by an LDAP server.
- User details can be specified in an XML file.
- You can specify local users on the appliance itself.

You can specify the following authorization methods:

- Access policies can be defined in an XML file.
- Access policies can be defined in local user groups.

You can map user groups looked up in an LDAP directory onto groups defined in an XML file or defined locally.

Where you have locally defined users, RBM can specify password policies and account policies for them. These policies define the rules governing password (such as minimum length, character types, and expiration periods) and those rules governing when accounts are locked out after failed log in attempts.

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 347](#) 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.

User authorization, credential mapping, and access profiles

The credential mapping part of role based management authorizes appliance users to use different features on the IBM MQ Appliance.

When you define credential mapping under role based management, you specify which resources appliance users have access to. Credential mapping provides a high degree of granularity in which resources can be excluded or included. For example, you can specify that a user can be mapped onto a set of credentials that allow modification of network settings, but prohibit changing user settings.

Access that appliance users have to resources is controlled by an access profile. The access profile defines the set of privileges for one or more resources on the appliance. Privileges for a resource can be one or more of the following permissions:

- Read
- Write
- Add
- Delete
- Execute

A bundle of access rights (also termed access policies) constitutes an access profile. An access profile can originate from either of the following credential mapping sources:

Local user group

Locally configured user group.

XML file

A file that defines access profiles.

After an appliance user is authenticated and the access profile is evaluated, the appliance enforces the established access profile. The IBM MQ Appliance web UI displays only resources that the user has access to, and the command line recognizes only commands for resources that the user has access to. For commands that users do not have access to, the command line displays the following message.

```
Unknown command or macro (command)
```

Access to IBM MQ resources

The following resources control administrative access by appliance users to IBM MQ on the appliance:

mq/cli

Granting execute permission on this resource allows the user to use the IBM MQ commands on the command line. The user can issue the **mqcli** command and administer the MQ aspects of the system. Users can access the command line by using SSH or by using the serial command line connection to the appliance. All permissions other than execute on the **mq/cli** resource are ignored.

mq/webadmin

Granting read/write permission to this resource allows the user to administer IBM MQ by using the IBM MQ Console and the MQ REST API. Granting read permission only allows the user only to view IBM MQ objects by using the IBM MQ Console and the MQ REST API, but not to change anything. All other permissions on this resource are ignored.

mq/webuser

This resource allows you to delegate the MQ authority checks for this appliance user to a matching messaging user. A messaging user with the same name must be defined on the appliance by using the messaging user and group commands (see [“Messaging user and group commands”](#) on page 624). The authorities for this messaging user are used for all operations in the IBM MQ Console and for the MQ REST API. Grant execute permission on the **mq/webuser** resource to represent ‘execution with delegated authority’. Other permissions against this resource are ignored.

mq/mftwebadmin

Granting read/write permission to this resource allows the user to perform all Managed File Transfer (MFT) operations using the MQ REST API. Granting read permission allows the user only to perform read only operations (GET requests), such as list transfer and list agents. All other permissions on this resource are ignored.

See [Roles on the IBM MQ Console and REST API](#) in the IBM MQ documentation for more details about IBM MQ resources.

Access policies

You use access policies to control which appliance resources users can access.

Access policies are strings that identify a particular resource and grant access to it. A number of access policies form an access profile, which can be applied to a particular user either through credential mapping using a local user group, or an XML file.

The access policy for the IBM MQ Appliance has the following format:

```
*/*/resource?Access=privileges
```

Where:

- *resource* is a URI that identifies the resource.
- *privileges* define the access given to the resource. Specify one or more of the following privileges, optionally separated by the plus (+) character:
 - r - read
 - w - write
 - x - execute
 - a - add
 - d - delete

You can also enter NONE to explicitly exclude users from a resource.

The following strings are examples of access policies:

```
*/**/?Access=r+w  
*/*/access/change-password?Access=x
```

A user with the access profile defined by these policies has read and write access to all appliance resources, plus they have execute permission on the `access/change-password` resource, which enables them to change their own password on the appliance.

There can be multiple matches when resolving access policies, and some of these might conflict with each other. In such cases, the more resource-specific policies are granted greater weight and override the more general policies. For example, a user group might have the following policies defined:

```
*/**/?Access=rwadx  
*/*/mgmt/rest-mgmt?Access=NONE
```

The first policy grants users read/write/add/delete/execute permissions to all appliance resources, but the second policy restricts access to the REST management interface such that no user in the group can access it. You can use these weighting characteristics to give general permissions to all appliance resources and add more policies that exclude specific resources.

Policies applying to the same users and resource that have the same weight are additive. So, for example, if a policy provides a user group with read access on a resource, and another policy provides write access on that resource, then users in that group have both read and write permissions on that resource.

When defining an access policy for a local user group, you can enter the profile strings manually using the `access-policy` command, or in the IBM MQ Appliance web UI. You can also use the policy builder in the web UI to specify access policies.

The appliance resources are listed in the following tables. The tables provide the following information for each resource:

- **Resource category.** The category the resource is listed under in the IBM MQ Appliance web UI.
- **Resource.** The name of the resource.
- **Resource profile URI.** The URI that you specify in an access policy giving access to this resource.
- **CLI command.** If you grant access to this resource, then users have access to this CLI command.

- **REST URI.** If you grant access to this resource, then users can use this REST URI (provided that they have access to the REST management interface).

Status resources

The status resources control access to status reporting for various aspects of appliance operation. Giving a user read access to a status resource enables them to use the **show** CLI command for that resource, or to use a REST query to recover the status of that resource.

Table 38. Status resources

Resource category	Resource	Resource profile URI	CLI command	REST URI
Main	Active services	status/active-services	show services	/mgmt/status/default/ServicesStatus
Main	Active Users	status/active-users	show users	/mgmt/status/default/ActiveUsers
Main	Date and Time	status/date-time	show time	/mgmt/status/default/DateTimeStatus
Main	Logging Targets	status/logging-target	show logging status	/mgmt/status/default/LogTargetStatus
Main	Object Status	status/object-status	show op-state	/mgmt/status/default/ObjectStatus
Main	Services Memory Usage	status/memory-services	show services-memory	/mgmt/status/default/ServicesMemoryStatus2
Configuration	Domain Status	status/domain-status	show domains	/mgmt/status/default/DomainStatus
System	Failure Notification	status/failure-notification	show failure-notification-status	/mgmt/status/default/FailureNotificationStatus2
System	Filesystem Information	status/filessystem	show filesystem	/mgmt/status/default/FilesystemStatus
System	Firmware Information	status/firmware	show firmware	/mgmt/status/default/FirmwareStatus2
System	Firmware Information	status/firmware-deprecated	-	/mgmt/status/default/FirmwareStatus
System	Version Information	status/firmware-version	show firmware-version	/mgmt/status/default/FirmwareVersion3
System	Version Information	status/firmware-version-deprecated	-	/mgmt/status/default/FirmwareVersion2
System	Hypervisor Information	status/hypervisor	show hypervisor	/mgmt/status/default/Hypervisor2
System	IPMI SEL Events	status/ipmi-sel-events	show ipmi-sel-events	/mgmt/status/default/IPMISelEvents

Table 38. Status resources (continued)

Resource category	Resource	Resource profile URI	CLI command	REST URI
System	Device Features	status/license	show features	/mgmt/status/default/LicenseStatus
System	Other Sensors	status/other-sensors	show sensors-other	/mgmt/status/default/OtherSensors
System	PCI Bus	status/pcibus	show pci-bus	/mgmt/status/default/PCIBus
System	RAID Array Status	status/raid-array	show raid-array	/mgmt/status/default/RaidArrayStatus
System	RAID Battery Backup Unit Status	status/raid-battery-module	show raid-battery-module	/mgmt/status/default/RaidBatteryModuleStatus
System	RAID Logical Drive Status	status/raid-logical-drive	show raid-logical-drive	/mgmt/status/default/RaidLogicalDriveStatus
System	RAID Physical Drive Status	status/raid-physical-drive	show raid-physical-drive	/mgmt/status/default/RaidPhysicalDriveStatus
System	SSD Status Information	status/raid-ssd	show raid-ssd	/mgmt/status/default/RaidSsdStatus
System	Fan Sensors	status/sensors-fans	show sensors-fans	/mgmt/status/default/EnvironmentalFanSensors
System	CPU usage	status/system-cpu	show system-cpu	/mgmt/status/default/SystemCpuStatus
System	Memory usage	status/system-memory	show system-memory	/mgmt/status/default/SystemMemoryStatus
System	Temperature Sensors	status/temperature	show sensors-temperature	/mgmt/status/default/TemperatureSensors
System	Virtual Appliance Information	status/virtual-platform	show virtual-platform	/mgmt/status/default/VirtualPlatform2
System	Voltage Sensors	status/voltage	show sensors-voltage	/mgmt/status/default/VoltageSensors
System	Current Sensors	status/current	show sensors-current	/mgmt/status/default/CurrentSensors

Table 38. Status resources (continued)

Resource category	Resource	Resource profile URI	CLI command	REST URI
IP-Network	Link Aggregation Status	status/agg	show link-aggregation-status	/mgmt/status/default/LinkAggregationStatus
IP-Network	DNS Cached Hosts	status/dns-cache-host	show dns-cache	/mgmt/status/default/DNSCacheHostStatus4
IP-Network	DNS Servers	status/dns-name-server	show name-servers	/mgmt/status/default/DNSNameServerStatus2
IP-Network	DNS Search Domains	status/dns-search-domain	show search-domains	/mgmt/status/default/DNSSearchDomainStatus
IP-Network	DNS Static Hosts	status/dns-static-host	show static-hosts	/mgmt/status/default/DNSStaticHostStatus
IP-Network	Ethernet Counters	status/ethernet-counters	show ethernet-counters	/mgmt/status/default/EthernetCountersStatus
IP-Network	Ethernet MAU	status/ethernet-mau	show ethernet-mau	/mgmt/status/default/EthernetMAUStatus
IP-Network	Ethernet MII Registers	status/ethernet-mii-registers	show ethernet-mii-registers	/mgmt/status/default/EthernetMIIRegisterStatus
IP-Network	IGMP Status	status/igmp-table	show igmp-table	/mgmt/status/default/IGMPStatus
IP-Network	IP address status	status/ipaddress	show ipaddress	/mgmt/status/default/IPAddressStatus
IP-Network	Link status	status/link	show link	/mgmt/status/default/LinkStatus
IP-Network	Link Aggregation Member Status	status/link-aggregation-member	show link-aggregation-member-status	/mgmt/status/default/LinkAggregationMemberStatus
IP-Network	Load Balancer Status	status/loadbalancer-status	show loadbalancer-status	/mgmt/status/default/LoadBalancerStatus2
IP-Network	ND Cache Table	status/ND-cache2	show ndcache	/mgmt/status/default/NDCacheStatus2
IP-Network	Network Interfaces	status/networkinterface	show network-interface	/mgmt/status/default/NetworkInterfaceStatus
IP-Network	Port Status	status/port-status	-	/mgmt/status/default/

Table 38. Status resources (continued)

Resource category	Resource	Resource profile URI	CLI command	REST URI
IP-Network	Routing Table	status/routing3	show route	/mgmt/status/default/RoutingStatus3
IP-Network	TCP Port Summary	status/tcp-connections	show tcp-connections	/mgmt/status/default/TCPSummary
IP-Network	TCP Port Status	status/tcp-table	show tcp-table	/mgmt/status/default/TCPTable
IP-Network	VLAN Interface Status	status/vlan2	show vlan-status	/mgmt/status/default/VlanInterfaceStatus2
Other Network	NFS Mount Status	status/nfs-mount	show nfs-mount	/mgmt/status/default/NFSMountStatus
Other Network	NTP Refresh Status	status/ntp-refresh	show ntp-refresh	/mgmt/status/default/NTPRefreshStatus
Other Network	SNMP Status	status/snmp-status	show snmp-status	/mgmt/status/default/SNMPStatus
  Other Network	Fibre Channel Host Bus Adapter	status/fibre-channel-hba	show fibre-channel-hba-status	/mgmt/status/default/FibreChannelHBAStatus
  Other Network	Discovered Fibre Channel LUNs	status/fibre-channel-luns	show fibre-channel-luns	/mgmt/status/default/FibreChannelLuns
  Other Network	Fibre Channel Volumes	status/fibre-channel-volume-status	show fibre-channel-volume-status	/mgmt/status/default/FibreChannelVolumeStatus
Crypto	Cryptographic Mode Status	status/crypto-mode	show crypto-mode	/mgmt/status/default/CryptoModeStatus
Crypto	SSH Known Host Table	status/trusted-hosts	show known-hosts	/mgmt/status/default/SSHTrustedHostStatus
MQ	MQ System Resources	status/mq-resources	show mq-resources	/mgmt/status/default/MQSystemResources
MQ	Queue Managers Status	status/qm-status	show qm-status	/mgmt/status/default/QueueManagersStatus

Configuration resources

The configuration resources give access to those resources that are used to configure the appliance. Giving a user permissions (read, write, add, and delete as required) to a configuration resource enables

them to work with configuration objects, using the web UI, or the CLI commands or REST URIs as listed in the following table.

<i>Table 39. Configuration resources</i>				
Resource category	Resource	Resource profile URI	CLI command	REST URI
Network Settings	DNS Settings	network/dns	config/dns	/mgmt/config/default/DNSNameService
Network Settings	Host Alias	network/host-alias	config/host-alias	/mgmt/config/default/HostAlias
Network Settings	Ethernet Interface	network/interface	config/ethernet	/mgmt/config/default/EthernetInterface
Network Settings	Link Aggregation Interface	network/link-aggregation	config/link-aggregation	/mgmt/config/default/LinkAggregation
Network Settings	Load Balancer Group	network/loadbalancer-group	config/loadbalancer-group	/mgmt/config/default/LoadBalancerGroup
Network Settings	Network Settings	network/network	config/network	/mgmt/config/default/NetworkSettings
Network Settings	NFS Client Settings	network/nfs-client	config/nfs-client	/mgmt/config/default/NFSClientSettings
Network Settings	NFS Static Mounts	network/nfs-static-mount	config/nfs-static-mount	/mgmt/config/default/NFSStaticMount
Network Settings	NTP Service	network/ntp-service	config/ntp-service config/ntp [deprecated]	/mgmt/config/default/NTPService
Network Settings	VLAN Interface	network/vlan	config/vlan	/mgmt/config/default/VLANInterface
  Network Settings	Fibre Channel Volume	network/fibre-channel-volume	config/fibre-channel-volume	/mgmt/config/default/FibreChannelVolume
Service Configuration	License Agent	services/ilmt-agent	config/ilmt-agent	/mgmt/config/default/ILMAgent
Crypto Configuration	Crypto Certificate	crypto/cert	config/crypto/certificate	/mgmt/config/default/CryptoCertificate
Crypto Configuration	Crypto Certificate Monitor	crypto/cert-monitor	config/crypto/cert-monitor	/mgmt/config/default/CertMonitor
Crypto Configuration	CRL Retrieval	crypto/crl	config/crypto/crl	/mgmt/config/default/CRLFetch
Crypto Configuration	Crypto Identification Credentials	crypto/idcred	config/crypto/idcred	/mgmt/config/default/CryptoIdentCred

Table 39. Configuration resources (continued)

Resource category	Resource	Resource profile URI	CLI command	REST URI
Crypto Configuration	Crypto Key	crypto/key	config/ crypto/key	/mgmt/config/default/ CryptoKey
Crypto Configuration	SSH Server Profile	crypto/ sshserverprofile	config/crypto/ sshserverprofile	/mgmt/config/default/ SSHServerProfile
Crypto Configuration	Crypto Shared Secret Key	crypto/sskey	config/crypto/ sskey	/mgmt/config/default/ CryptoSSKey
Crypto Configuration	SSL Client Profile	crypto/ssl-client	config/crypto/ ssl-client	/mgmt/config/default/ SSLClientProfile
Crypto Configuration	SSL Server Profile	crypto/ssl-server	config/crypto/ ssl-server	/mgmt/config/default/ SSLServerProfile
Crypto Configuration	SSL Host Name Mapping	crypto/ssl-sni-mapping	config/crypto/ ssl-sni-mapping	/mgmt/config/default/ SSLSNIMapping
Crypto Configuration	SSL SNI Server Profile	crypto/ssl-sni-server	config/crypto/ ssl-sni-server	/mgmt/config/default/ SSLSNIServerProfile
Crypto Configuration	Test Password Map	crypto/test-password-map	config/crypto/ test password-map	/mgmt/config/default/ TestPasswordMap
Crypto Configuration	Crypto Validation Credentials	crypto/valcred	config/crypto/ valcred	/mgmt/config/default/ CryptoValCred
Device Management	IPMI LAN Channel	mgmt/ipmi-lan-channel	config/ipmi-lan-channel	/mgmt/config/default/ IPMILanChannel
Device Management	IPMI User	mgmt/ipmi-user	config/ipmi-user	/mgmt/config/default/IPMIUser
Device Management	REST Management Interface	mgmt/rest-mgmt	config/rest-mgmt	/mgmt/config/default/ RestMgmtInterface
Device Management	SSH Service	mgmt/ssh	config/ssh	/mgmt/config/default/ SSHService
Device Management	Web Management Service	mgmt/web-mgmt	config/web-mgmt config/save-config overwrite	/mgmt/config/default/WebGUI
Access Settings	Access Control List	access/acl	config/acl	/mgmt/config/default/ AccessControllist
Access Settings	LDAP Search Parameters	access/ldap-search-parameters	config/ldap-search-parameters	/mgmt/config/default/ LDAPSearchParameters
Access Settings	RBM Settings	access/rbm	config/rbm	/mgmt/config/default/ RBMSettings
Access Settings	SNMP Settings	access/snmp	config/snmp	/mgmt/config/default/ SNMPSettings

Table 39. Configuration resources (continued)

Resource category	Resource	Resource profile URI	CLI command	REST URI
Access Settings	User Group	access/usergroup	config/usergroup	/mgmt/config/default/UserGroup
Access Settings	User Account	access/username	config/user	/mgmt/config/default/User
Configuration Management	Password Map Alias	config/password-alias	config/password-alias	/mgmt/config/default/PasswordAlias
Configuration Management	Password Map	config/password-map	-	-
Logging Configuration	Audit Log Settings	logging/audit-log	config/audit-log-settings	/mgmt/config/default/AuditLogSettings
Logging Configuration	Log Category	logging/category	config/loggingcategory	/mgmt/config/default/LogLabel
Logging Configuration	Log Target	logging/target	config/loggingtarget	/mgmt/config/default/LogTarget
System Settings	Failure Notification	system/failure-notification	config/failure-notification	/mgmt/config/default/ErrorReportSettings
System Settings	Language	system/language	config/language	/mgmt/config/default/Language
System Settings	RAID Array	system/raid-disk-volume	config/raid-volume	-
System Settings	System Settings	system/system	config/system config/globallogipfilter	/mgmt/config/default/SystemSettings
System Settings	Time Settings	system/timezone	config/timezone	/mgmt/config/default/TimeSettings

Action resources

The action resources control access to the resources used to perform actions on the appliance. Give users execute permission on a resource to enable the corresponding action. Users can perform the action by using the corresponding CLI command or by sending a request to the REST URI. All action requests use the URI /mgmt/actionqueue/default/operations. The REST column in the following table gives the operation name used when constructing a payload to request an action (see “[Triggering appliance operations by using the REST management interface](#)” on page 370).

Table 40. Action resources

Resource category	Resource	Resource profile URI	CLI command	Operation name for REST request
Device Settings	Add IPMI BMC SEL Test Entry	device/add-ipmi-sel-test-entry	config/add-ipmi-sel-test-entry	AddSelTestEntry
Device Settings	Delete previous firmware install	device/boot-delete	config/flash/bootdelete	BootDelete
Device Settings	Boot Image	device/boot-image	config/flash/bootimage	ApplyPatch

Table 40. Action resources (continued)

Resource category	Resource	Resource profile URI	CLI command	Operation name for REST request
Device Settings	Switch Install Image	device/boot-switch	config/flash/boot-switch	BootSwitch
Device Settings	Boot Update	device/boot-update	config/flash/boot-update	BootUpdate
Device Settings	Clear IPMI BMC SEL	device/clear-ipmi-sel	config/clear-ipmi-sel	ClearSel
Device Settings	Create Directory	device/create-dir	config/mkdir	CreateDir
Device Settings	Delete File	device/delete-file	config/delete	DeleteFile
Device Settings	Fetch File	device/fetch-file	config/copy	FetchFile
Device Settings	Initialize file system	device/initialize-raid-volume-filesystem	config/raid-volume-initialize-filesystem	-
Device Settings	Control Locate LED	device/locate-device	config/locate-device	-
Device Settings	Move File	device/move-file	config/move	MoveFile
Device Settings	Activate RAID Array	device/raid-activate	config/raid-activate	-
Device Settings	Delete RAID Array	device/raid-delete	config/raid-delete	-
Device Settings	Initialize RAID Array	device/raid-initialize	config/raid-initialize	-
Device Settings	Request Learning Cycle for BBU	device/raid-learn-battery	config/raid-learn-battery	-
Device Settings	Make hot spare for RAID Array	device/raid-make-hot-spare	config/raid-make-hot-spare	-
Device Settings	Rebuild RAID Array	device/raid-rebuild	config/raid-rebuild	-
Device Settings	Remove Directory	device/remove-dir	config/rmdir	RemoveDir
Device Settings	Send File	device/sendfile	config/send file	SendFile
Device Settings	Shut down	device/shutdown	config/shutdown	Shutdown
Device Settings	Set Time and Date	device/time-date	config/clock	SetTimeAndDate
Device Settings	VerifyFirmware	device/verify-firmware	config/flash/verify-firmware	VerifyFirmware
Device Settings	Control Locate LED	device/locate-device	config/locate-device	LocateDevice
Network Settings	Quiesce	network/quiesce	-	-

Table 40. Action resources (continued)

Resource category	Resource	Resource profile URI	CLI command	Operation name for REST request
Network Settings	Unquiesce	network/unquiesce	-	-
  Network Settings	Initialize Filesystem on Fibre Channel Device	network/fibre-channel-fs-init	config/fibre-channel-fs-init	InitFibreChannelFilesystem
  Network Settings	Repair Filesystem on Fibre Channel Volume	network/fibre-channel-fs-repair	config/fibre-channel-fs-repair	RepairFibreChannelFilesystem
  Network Settings	Unlock Filesystem on Fibre Channel Volume	network/fibre-channel-unlock-volume	config/fibre-channel-unlock-volume	UnlockFibreChannelVolume
Crypto Configuration	Add Password Map	crypto/add-password-map	config/crypto/password-map	AddPasswordMap
Crypto Configuration	Convert Crypto Certificate Object	crypto/convert-certificate	config/crypto/convert-certificate	ConvertCertificate
Crypto Configuration	Convert Crypto Key Object	crypto/convert-key	config/crypto/convert-key	ConvertKey
Crypto Configuration	Export Crypto Object	crypto/crypto-export	config/crypto/crypto-export	CryptoExport
Crypto Configuration	Import Crypto Object	crypto/crypto-import	config/crypto/crypto-import	CryptoImport
Crypto Configuration	Set Cryptographic Mode	crypto/crypto-mode-set	config/crypto/crypto-mode-set	CryptoModeSet
Crypto Configuration	Delete Password Map	crypto/delete-password-map	config/crypto/delete password-map	DeletePasswordMap
Crypto Configuration	Generate Key	crypto/keygen	config/crypto/keygen	Keygen
Crypto Configuration	Delete SSH Known Host	crypto/no-known-host	config/crypto/no-client-known-host	DeleteKnownHost
Crypto Configuration	Delete SSH Known Host Table	crypto/no-known-host-table	config/crypto/no-client-known-host-table	DeleteKnownHostTable
Crypto Configuration	No Password Map	crypto/no-password-map	config/crypto/no-password-map	NoPasswordMap

Table 40. Action resources (continued)

Resource category	Resource	Resource profile URI	CLI command	Operation name for REST request
Crypto Configuration	test-password-map	crypto/test-password-map	config/crypto/test password-map	TestPasswordMap
Access Settings	Change User Password	access/change-password	config/user-password	ChangePassword
Access Settings	Disconnect	access/disconnect	config/disconnect	Disconnect
Access Settings	Force Password Change	access/force-password-change	config/user-expire-password	UserForcePasswordChange
Access Settings	Reset Failed Login Counter	access/reset-failed-login	config/reset-failed-login	UserResetFailedLogin
Access Settings	Reset Password	access/reset-username	config/reset-username	UserResetPassword
Configuration Management	Execute Configuration	config/exec-config	config/exec	ExecConfig
Configuration Management	Password Map	config/password-map	-	-
Configuration Management	REST Export	config/rmi-export	-	Export
Configuration Management	REST Load Configuration	config/rmi-load-config	-	LoadConfiguration
Configuration Management	View Certificate Details	config/rmi-view-details	-	ViewCertificateDetails
Configuration Management	Save Configuration	config/save-config	config/write-memory	SaveConfig
Configuration Management	Save Internal State	config/saveinternlstate	config/save-internal-state	SaveInternalState
Configuration Management	Select Configuration	config/select-config	config/flash/boot-config	SelectConfig
Configuration Management	Undo Configuration	config/undo-config	config/undo	UndoConfig
System Settings	Delete SSH Known Host	system/no-trusted-host	config/no-known-host	DeleteTrustedHost
System Settings	Add SSH Known Host	system/trusted-host	config/known-host	AddTrustedHost
Cache Management	Flush ARP Cache	cache/flush-arp	config/clear arp	FlushArpCache
Cache Management	Flush DNS Cache	cache/flush-dns	config/clear dns-cache	FlushDNSCache
Cache Management	Flush ND Cache	cache/flush-ndcache	config/clear-ndcache	FlushNDCache

Table 40. Action resources (continued)

Resource category	Resource	Resource profile URI	CLI command	Operation name for REST request
Cache Management	Flush RBM Cache	cache/flush-rbm	config/clear rbm cache	FlushRBMCache
Cache Management	Flush Document	cache/refresh-document	-	RefreshDocument
Debug Settings	Disable hardware offload	debug/disable-aggregation-hardware-offload	config/disable-aggregation-hardware-offload	DisableLinkAggregationHardwareOffload
Debug Settings	Disable hardware offload	debug/disable-ethernet-hardware-offload	config/disable-ethernet-hardware-offload	DisableEthernetHardwareOffload
Debug Settings	Disable hardware offload	debug/disable-vlan-hardware-offload	config/disable-vlan-hardware-offload	DisableVLANHardwareOffload
Debug Settings	Generate Error Report	debug/error-report	config/save error-report	ErrorReport
Debug Settings	Start packet capture	debug/packet-capture	config/ethernet <name>/packet-capture config/link-aggregation <name>/packet-capture config/packet-capture-advanced config/vlan <name>/packet-capture	DisableLinkAggregationHardwareOffload DisableEthernetHardwareOffload DisableVLANHardwareOffload ErrorReport PacketCapture PacketCaptureDebug StopPacketCapture LinkAggregationPacketCapture LinkAggregationStopPacketCapture UniversalPacketCaptureDebug UniversalStopPacketCapture VLANPacketCapture VLANStopPacketCapture
Debug Settings	Ping Remote	debug/ping	config/ping	Ping
Debug Settings	Send Error Report	debug/send-error-report	config/send error-report	SendErrorReport
Debug Settings	Generate Log Event	debug/send-logevent	config/test logging	SendLogEvent
Debug Settings	Set Log Level	debug/set-loglevel	config/loglevel	SetLogLevel
Debug Settings	Enable RBM Debug Logging	debug/set-rbmlog	-	SetRBMDebugLog
Debug Settings	TCP Connection Test	debug/tcp-connection-test	config/test tcp-connection	TCPConnectionTest
Debug Settings	Hardware Diagnostics	debug/test-hardware	config/test hardware	TestHardware

Table 40. Action resources (continued)

Resource category	Resource	Resource profile URI	CLI command	Operation name for REST request
Debug Settings	Disable Ethernet Hardware Offload	debug/disable-ethernet-hardware-offload	config/disable-ethernet-hardware-offload	DisableEthernetHardwareOffload
Debug Settings	Disable Link Aggregation Hardware Offload	debug/disable-aggregation-hardware-offload	config/disable-aggregation-hardware-offload	DisableLinkAggregationHardwareOffload
Debug Settings	Disable VLAN Hardware Offload	debug/disable-vlan-hardware-offload	config/disable-vlan-hardware-offload	DisableVLANHardwareOffload

Admin only resources

The resources listed in the following table are only visible to, and usable by, the admin user. You cannot alter access to these resources.

Table 41. Admin only resources

Resource	Resource profile URI	CLI command
Diagnostics	Only available to admin user	diagnostics
Trace Route	Only available to admin user	traceroute
Clear Intrusion Detected	Only available to admin user	clear intrusion-detected
Watchdog	Only available to admin user	config/watchdog
Startup Configuration	Only available to admin user	config/startup
Reinitialize	Only available to admin user	config/flash/reinitialize
Service Nagle	Only available to admin user	config/service nagle
Log Size	Only available to admin user	config/logsize
System Log	Only available to admin user	config/syslog

Other resources

The following table lists the resources in the following groups:

- Login - permissions on these resources specify which interfaces users can use to interact with the appliance. There are no CLI commands or REST URIs associated with these resources.
- File management - permissions on these resources give users access to directories on the appliance.
- MQ configuration - permissions on these resources give users access to IBM MQ on the appliance.

The following CLI commands are always available to all users who can connect to the command line:

- echo
- exit
- help
- login
- top
- template

- config/dir

Table 42. Other resources

Resource category	Resource	Resource profile URI	CLI command	REST URI
Login	SSH	login/ssh	-	-
Login	Web-Mgmt	login/web-mgmt	-	-
Login	Rest-Mgmt	login/rest-mgmt	-	-
File Management	local:	file/local	-	/mgmt/filestore/default/local
File Management	temporary:	file/temporary	-	/mgmt/filestore/default/temporary
File Management	store:	file/store	-	/mgmt/filestore/default/store
File Management	config:	file/config	-	/mgmt/filestore/default/config
File Management	image:	file/image	-	/mgmt/filestore/default/image
File Management	logstore:	file/logstore	-	/mgmt/filestore/default/logstore
File Management	logtemp:	file/logtemp	-	/mgmt/filestore/default/logtemp
File Management	audit:	file/audit	-	/mgmt/filestore/default/audit
File Management	tasktemplates:	file/ tasktemplates	-	/mgmt/filestore/default/tasktemplates
File Management	cert:	file/cert	-	/mgmt/filestore/default/cert
File Management	pubcert:	file/pubcert	-	/mgmt/filestore/default/pubcert
File Management	sharedcert:	file/sharedcert	-	/mgmt/filestore/default/sharedcert
File Management	export:	file/export	-	/mgmt/filestore/default/export
File Management	mqbackup:	file/mqbackup		/mgmt/filestore/default/mqbackup
File Management	mqdiag:	file/mqdiag		/mgmt/filestore/default/mqdiag
File Management	mqerr:	file/mqerr		/mgmt/filestore/default/mqerr
File Management	mqpubcert:	file/mqpubcert		/mgmt/filestore/default/mqpubcert
File Management	mqmqdata:	file/mqqmdata		/mgmt/filestore/default/mqqmdata
File Management	mqtemporary:	file/ mqtemporary		/mgmt/filestore/default/mqtemporary
File Management	mqtrace:	file/mqtrace		/mgmt/filestore/default/mqtrace
File Management	mqwebui:	file/mqwebui		/mgmt/filestore/default/mqwebui
File Management	fcvolumes:	file/fcvolumes		/mgmt/filestore/default/fcvolumes
MQ Configuration	MQ CLI Administration	mq/cli	mqcli	-
MQ Configuration	MQ Web Administration	mq/webadmin	-	-
MQ Configuration	MQ Web User	mq/webuser	-	-
MQ Configuration	MQ MFT Administration	mq/ mftwebadmin	-	-

Role based management

Appliance users and their permissions are controlled by role based management.

You configure role based management to determine how users logging into the appliance are authenticated. You also set up access profiles to determine what appliance resources users can work with after they are authenticated.

You can configure role based management by using the IBM MQ Appliance web UI or by using the command line interface.

If you use external LDAP servers for user authentication, be aware that these servers might potentially be a weakness in your security setup. You must take the necessary steps to ensure that the LDAP servers are themselves secure.

User authentication

You can configure role based management to authenticate users in one of the following ways:

LDAP

The appliance authenticates users remotely by using an LDAP server. You can also define local users to fall back to if the LDAP server is not available.

Local user

When authentication is local, authentication is performed by the appliance by using user name and password.

XML file

User names and passwords can be specified in an XML file. You can store the XML file on the appliance or on a remote server. You can use the RBM builder on the appliance to define users. You can use the same XML file to define access policies.

User authorization

You can configure role based management to authorize users to use appliance resources by selecting one of the following credential mapping methods:

Local user group

Specify access profiles in the local user groups on the appliance. You can map user groups or individual users looked up on an LDAP server onto local user groups, which allows a user to belong to multiple role-based groups.

XML file

Specify access policies in an XML file. You can store the XML file on the appliance or on a remote server. You can use the RBM builder on the appliance to define access profiles. You can map user groups or individual users looked up on an LDAP server onto policies that are defined in an XML file.

User authorization enforces access privileges for one or more resources on the appliance. These privileges can be quite broad or very specific. The privileges are combined together to form an access profile. See [“User authorization, credential mapping, and access profiles” on page 423](#) for detailed information.

The following table illustrates the permitted mixes of authentication and authorization methods on the appliance.

	Local user group authorization	XML file authorization
LDAP authentication	Yes	Yes
Local user authentication	Yes	Yes
XML file authentication	Yes	Yes

Important: avoiding user lock out when configuring role based management

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

If your user authentication depends on one or more external LDAP servers, for example, then you must take steps to ensure that log in is still possible for one or more users if you lose connection to the external servers. You do this by configuring one or more fallback users. Fallback users are local users who are authenticated by the appliance.

The RBM settings that you configure apply to users that are accessing the appliance both by the web UI and by the CLI. (If you do lock yourself out, you can attach a terminal directly to the physical appliance and log in as user `admin`.)

When you configure user authentication by using the web UI, changes take effect as soon as you click **Apply**. Be careful that you do not lock yourself out before you have defined fallback users, and ensure that your authentication server is available and appropriately configured.

You can use the following steps to double-check your changes (you might require physical access to your appliance in order to restart it):

1. Complete the required RBM modifications, including the definition of one or more fallback users (and an LDAP load balancer group, if you are using one).
2. Click **Apply** to enforce the changes, but do not click **Save Config**.
3. Verify that one of your fallback users can log in to the appliance.
4. Verify that one of your externally verified users can log in to the appliance.
5. If your externally verified user cannot log in, make the necessary changes as the fallback user, apply your changes, and try again.
6. If your fallback user cannot log in, then physically restart the appliance (by using the power button) to roll back to the previously saved configuration.

User authentication with LDAP

You can configure the IBM MQ Appliance to authenticate users by using an LDAP server.



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management”](#) on page 440 before you attempt to set up authentication with LDAP

You have a number of options when you are using an LDAP server:

Using credentials directly

You can use your users' credentials directly to bind to the LDAP server. In this case, your user names must be part of the X.500 distinguished names (DN) that the LDAP server uses to identify directory entries. You specify the remainder of the DN as part of the configuration. Typically, the user would log in using the common name (CN) part of the distinguished name. The appliance prefixes the user name with "cn=" and suffixes it with a comma and the remaining distinguished name elements that are common to all appliance users.

For example, your user might log in with the user name "Robin Dalemain", which is the CN part of their DN. You have configured the suffix to be "dc=appliance203, dc=com". When Robin attempts to log in to the appliance, the distinguished name "cn=Robin Dalemain, dc=appliance203, dc=com", together with Robin's password are used to connect to the LDAP server. If Robin's credentials successfully bind to the LDAP server, then Robin is authenticated and can access the appliance.

Looking up users in LDAP

You can configure the appliance so that users enter a user name that is not part of their DN. You specify search parameters so that this user name can be passed to the LDAP server and used to look up the user's distinguished name, which is then used with the entered password to authenticate the user. To look up a user's distinguished name the appliance can either bind to the LDAP server anonymously, or you can specify a bind ID and password alias it must use.

For example, Robin Dalemain might have the user name "RWD123". When Robin attempts to log in to the appliance, Robin's user name is sent to the LDAP server and the distinguished name for Robin's entry is returned. Robin is authenticated using his distinguished name and password to determine if he can access the appliance.

Using TLS (SSL)

You can specify that the appliance acts as an SSL client when connecting to the LDAP server. If you use this option, user credentials are encrypted when sent to the LDAP server, so user passwords are never sent across the network in plain text.

Using load balancing

You can specify that the appliance uses a pool of LDAP servers rather than a single server, and configure how the load is balanced between the LDAP servers in the pool.

Specifying fallback users

It is important that you specify one or more fallback users. These are local users who can log in to the appliance if you lose the connection with your LDAP server.

After you have configured how users are authenticated using LDAP, you must go on to specify how authenticated users are authorized to use the appliance resources. You do this by configuring credential mapping.

Related concepts

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Configuring direct authentication with LDAP by using the web UI

Configure the appliance to pass user credentials to the LDAP server and use them as the bind credentials.

Before you begin



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management” on page 440](#) before you attempt to set up authentication with LDAP

About this task

You can use the IBM MQ Appliance web UI to configure role based management such that the appliance uses user credentials directly to bind the LDAP server. If the bind is successful, the user who is attempting to log in to the appliance is successfully authenticated. See [“User authentication with LDAP” on page 440](#) for a description of this method of authentication.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > RBM Settings**
3. Ensure that **Enable administrative state** is selected (it is selected by default) and click **Authentication** to view the authentication options.
4. Select an **Authentication method of LDAP**.
5. Specify the **Server host** and the **Server port** for connecting to the LDAP server (server port is usually 389, or 636 for an SSL connection), and select the **LDAP version** (the version is usually v3).
6. If you have configured a load balancer group for LDAP access and created a profile, specify it in the **Load balancer group** field. Alternatively, click the plus icon  to open the Load Balancer Group dialog to specify a profile for your load balancer group (see [“Creating a load balancer group profile by using the web UI” on page 451](#)). You can leave this field blank if you are using a single LDAP server.
7. Specify the **LDAP prefix** that the appliance prefixes the user name with when it is constructing a DN to pass to the LDAP server. The prefix is cn= by default.

- Specify the **LDAP suffix** that the appliance appends to the user name when it is constructing a DN to pass to the LDAP server. For example, "dc=appliance123,dc=com".
- Specify an **LDAP read timeout**. The timeout is the time that the appliance will attempt to connect to the LDAP server before closing the connection. The default is 60 seconds. Specify 0 to never timeout.
- If you want to use an SSL (TLS) connection to the LDAP server, select an **SSL client type** of **Client profile**. If you have already defined a profile, select the profile name from the **SSL client profile** list.

Alternatively, click the plus icon  to open the SSL Client Profile dialog and create a new SSL client profile (see [“Creating an SSL client profile by using the web UI”](#) on page 448.)

- To define fallback users, you can choose **All users** from the **Local accounts for fallback** list to have all local users able to log in to the appliance if LDAP is unavailable. Alternatively, select **Specific users** and select one or more local users.
- Optionally, change the default cache settings. Cache settings determine how long user details are held on the appliance before authentication is referred to the LDAP server again. By default, the appliance retains details for an absolute period of 600 seconds. You can change the cache mode or the cache lifetime, or both. You can also disable caching altogether.
- Click **Apply** to apply your changes.

What to do next

After you specify the user authentication method, you must next configure credential mapping.

Related concepts

[“Important: avoiding user lock out when configuring role based management”](#) on page 440

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Configuring direct authentication with LDAP by using the command line

Use the command line to configure the appliance to pass user credentials to the LDAP server and use them as the bind credentials.

Before you begin



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management”](#) on page 440 before you attempt to set up authentication with LDAP

About this task

You can use commands to configure role based management such that the appliance uses user credentials directly to bind the LDAP server. If the bind is successful, the user who is attempting to log in to the appliance is successfully authenticated. See [“User authentication with LDAP”](#) on page 440 for a description of this method of authentication.

Procedure

- Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207. Log in as an administrative user.
- Type `config` to enter global configuration mode.
- Type the following command to configure role based management:

```
rbm
```

- Enter the following command to specify the LDAP authentication method:

```
au-method ldap
```

- Specify the host name of the LDAP server (you can use a host alias if you have defined one):

```
au-server-host host
```

6. Specify the port to use when connecting to the LDAP server:

```
au-server-port port
```

The usual port for an LDAP server is 389, or 636 for an SSL connection.

7. Specify the LDAP version that is used to access the LDAP server for RBM authentication:

```
ldap-version version
```

Where *version* is v2 or v3. The default value is v2.

8. Optionally specify a load balancer group. (See [“Creating a load balancer group profile by using the command line”](#) on page 453 for details of how to create a load balancer group):

```
loadbalancer-group name
```

9. Specify the string that is used to prefix the user name, for example, "cn=".

```
ldap-prefix prefix
```

10. Specify the string that is used to suffix user names, for example, "dc=appliance123, dc=com".

```
ldap-suffix suffix
```

11. Specify the time that RBM authentication waits for a response from the LDAP server. The default value is 60. A value of 0 indicates that the wait never times out:

```
au-ldap-readtimeout seconds
```

12. If you need a secure connection with your LDAP server, you must specify an SSL client type of client, and the name of your SSL client profile (see [“Creating an SSL client profile by using the command line”](#) on page 449 for details of how to create an SSL client profile):

```
ssl-client-type client  
ssl-client name
```

13. Optionally specify fallback users who can log in to the appliance if the LDAP server is not available. Fallback users must already have been added as local users to the appliance. You can specify that all local users are fallback users by entering the following command:

```
fallback-login local
```

Alternatively, you can specify one or more particular users by entering the following commands:

```
fallback-login restricted  
fallback-user localuser1  
fallback-user localuser2  
⋮  
fallback-user localuserN
```

14. Alter the default LDAP cache settings, if required. By default, the appliance caches results of authentication attempts for 600 seconds, but you can change the mode of caching, and the caching duration by entering the following commands:

```
au-cache-mode mode  
au-cache-ttl seconds
```

Where *mode* is one of:

absolute

Caches the results of user authentications for a period of time specified by the **au-cache-ttl** command (the explicit time-to-live). This is the default setting.

disabled

Disables caching. The appliance will not cache any results and instead always authenticates every time a user requests access.

maximum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the maximum of the two values.

minimum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the minimum of the two values.

Example

The following example configures the appliance to use an LDAP server that is identified by the host alias "ldap_host" for user authentication. If a user attempts to log in with the user name "Robin Dalemain", the string "cn=Robin Dalemain, dc=appliance123, dc=com" is passed to the LDAP server and used as the bind ID. Robin's password is used as the bind password. If the LDAP server is unavailable, any local appliance user can log in to the appliance.

```
mqa# config
Global configuration mode
mqa(config)# rbm
Modify RBM Settings configuration

mqa(config rbm)# au-method ldap
mqa(config rbm)# au-server-host ldap_host
mqa(config rbm)# au-server-port 389
mqa(config rbm)# ldap-version v3
mqa(config rbm)# ldap-prefix "cn="
mqa(config rbm)# ldap-suffix "dc=appliance123,dc=com"
mqa(config rbm)# fallback-login local
mqa(config rbm)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

What to do next

After you specify the user authentication method, you must next configure credential mapping.

Related concepts

[“RBM setting commands” on page 917](#)

Use the RBM commands to configure role based management on the IBM MQ Appliance.

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Configuring look up authentication with LDAP by using the web UI

Configure the appliance to look up user details on the LDAP server. To look up users, the appliance binds to the LDAP server by using credentials.

Before you begin



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management” on page 440](#) before you attempt to set up authentication with LDAP

About this task

You can use the IBM MQ Appliance web UI to configure role based management such that the appliance looks up user details in the LDAP server by using defined search parameters. To look up users, the appliance binds to the LDAP server by using credentials that you define as part of the RBM configuration,

or you can use an anonymous bind to access the LDAP server. See [“User authentication with LDAP” on page 440](#) for a description of this method of authentication.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > RBM Settings**
3. Ensure that **Enable administrative state** is selected (it is selected by default) and click **Authentication** to view the authentication options.
4. Select an **Authentication method** of **LDAP**.
5. Specify the **Server host** and the **Server port** for connecting to the LDAP server (server port is usually 389, or 636 for an SSL connection), and select the **LDAP version** (the version is usually v3).
6. If you have configured a load balancer group for LDAP access and created a profile, specify it in the **Load balancer group** field. Alternatively, click the plus icon  to open the Load Balancer Group dialog to specify a profile for your load balancer group (see [“Creating a load balancer group profile by using the web UI” on page 451](#)). You can leave this field blank if you are using a single LDAP server.
7. Select **Search LDAP for DN**.
8. Specify the DN that the appliance uses to bind to the LDAP server to perform the search in the **LDAP bind DN** field. Specify the password alias in the **LDAP bind password alias** field. Click the plus icon  to create a password alias if you have not already created one. (Leave these fields blank if you are using an anonymous bind to access the LDAP server.)
9. Specify the **LDAP search parameters**. You can enter these parameters directly, or you can click the plus icon to open the LDAP Search Parameters dialog.
10. Specify an **LDAP read timeout**. The timeout is the time that the appliance will attempt to connect to the LDAP server before closing the connection. The default is 60 seconds. Specify 0 to never timeout.
11. If you want to use an SSL (TLS) connection to the LDAP server, select an **SSL client type** of **Client profile**. If you have already defined a profile, select the profile name from the **SSL client profile** list.
Alternatively, click the plus icon  to open the SSL Client Profile dialog and create a new SSL client profile (see [“Creating an SSL client profile by using the web UI” on page 448](#)).
12. To define one or more fallback users, choose **All users** from the **Local accounts for fallback** list to have all local users able to log in to the appliance if LDAP is unavailable. Alternatively, select **Specific users** and select one or more local users.
13. Optionally, change the default cache settings. Cache settings determine how long user details are held on the appliance before authentication is referred to the LDAP server again. By default, the appliance retains details for an absolute period of 600 seconds. You can change the cache mode or the cache lifetime, or both. You can also disable caching altogether.
14. Click **Apply** to apply your changes.

What to do next

After you specify the user authentication method, you must next configure credential mapping.

Related concepts

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Configuring look up authentication with LDAP by using the command line

Use the command line to configure the appliance to look up user details on an LDAP server. To look up users, the appliance binds to the LDAP server by using credentials.

Before you begin



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management” on page 440](#) before you attempt to set up authentication with LDAP

About this task

You can use commands to configure role based management such that the appliance looks up user details in the LDAP server by using defined search parameters. To look up users, the appliance binds to the LDAP server by using credentials that you define as part of the RBM configuration, or you can use an anonymous bind to access the LDAP server. See [“User authentication with LDAP” on page 440](#) for a description of this method of authentication.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure role based management:

```
rbm
```

4. Enter the following command to specify the LDAP authentication method:

```
au-method ldap
```

5. Specify the host name of the LDAP server (you can use a host alias if you have defined one):

```
au-server-host host
```

6. Specify the port to use when connecting to the LDAP server:

```
au-server-port port
```

The usual port for an LDAP server is 389, or 636 for an SSL connection.

7. Specify the LDAP version that is used to access the LDAP server for RBM authentication:

```
ldap-version version
```

Where *version* is `v2` or `v3`. The default value is `v2`.

8. Optionally specify a load balancer group. (See [“Creating a load balancer group profile by using the command line” on page 453](#) for details of how to create a load balancer group):

```
loadbalancer-group name
```

9. Specify that the appliance will search the LDAP directory for user information:

```
au-ldap-search on
```

10. Specify the distinguished name used to bind to the LDAP server to perform the search:

```
au-ldap-bind-dn DN
```

11. Specify the password that is used for binding to the LDAP server. You should define a password alias, rather than entering the password directly (see [“password-map”](#) on page 740 for details on creating a password alias).

```
au-ldap-bind-password password_alias
```

If you do not specify a bind DN and bind password, an anonymous bind is used for the search.

12. Define the parameters for the search. You use the LDAP search parameters commands to create a named set of parameters (see [“LDAP Search Parameters commands”](#) on page 835). You then specify the name of the set of search parameters to use:

```
au-ldap-parameters name
```

13. Specify the time that RBM authentication waits for a response from the LDAP server. The default value is 60. A value of 0 indicates that the wait never times out:

```
au-ldap-readtimeout seconds
```

14. If you need a secure connection to your LDAP server, you must specify an SSL client type of client, and the name of your SSL client profile (see [“Creating an SSL client profile by using the command line”](#) on page 449 for details of how to create an SSL client profile):

```
ssl-client-type client  
ssl-client name
```

15. Optionally specify fallback users who can log in to the appliance if the LDAP server is not available. Fallback users must already have been added as local users to the appliance. You can specify that all local users are fallback users by entering the following command:

```
fallback-login local
```

Alternatively, you can specify one or more particular users by entering the following commands:

```
fallback-login restricted  
fallback-user localuser1  
fallback-user localuser2  
..  
fallback-user localuserN
```

16. Alter the default LDAP cache settings, if required. By default, the appliance caches results of authentication attempts for 600 seconds, but you can change the mode of caching, and the caching duration by entering the following commands:

```
au-cache-mode mode  
au-cache-ttl seconds
```

Where *mode* is one of:

absolute

Caches the results of user authentications for a period of time specified by the **au-cache-ttl** command (the explicit time-to-live). This is the default setting.

disabled

Disables caching. The appliance will not cache any results and instead always authenticates every time a user requests access.

maximum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the maximum of the two values.

minimum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the minimum of the two values.

Example

The following example configures the appliance to use an LDAP server that is identified by the host alias `ldap_host` for user authentication. If a user logged in with the user name "RobinWD", the name is passed to the LDAP server and used to look up Robin's distinguished name (as specified in the search parameters). If the LDAP server is unavailable, any local appliance user can log in to the appliance.

```
mqa# config
Global configuration mode
mqa(config)# rbm
Modify RBM Settings configuration

mqa(config rbm)# au-method ldap
mqa(config rbm)# au-server-host ldap_host
mqa(config rbm)# au-server-port 389
mqa(config rbm)# ldap-version v3
mqa(config rbm)# au-ldap-search on
mqa(config rbm)# au-ldap-bind-dn cn=appbind, dc=appliance123, dc=com
mqa(config rbm)# au-ldap-bind-password bindpw_alias
mqa(config rbm)# au-ldap-parameters auth_params
mqa(config rbm)# fallback-login local
mqa(config rbm)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

What to do next

After you specify the user authentication method, you must next configure credential mapping.

Related concepts

[“RBM setting commands” on page 917](#)

Use the RBM commands to configure role based management on the IBM MQ Appliance.

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Creating an SSL client profile by using the web UI

If you want to configure a secure SSL (TLS) connection with an LDAP server, then you must configure a client profile.

Before you begin



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management” on page 440](#) before you attempt to set up authentication with LDAP

About this task

You can configure an SSL client profile by using the IBM MQ Appliance web UI. You can do this by opening a dialog while you configure role based management (RBM), or as a separate operation before you configure RBM.

Procedure

1. Open the SSL Client Profile window. You can do this in one of two ways:

Before you configure RBM:

- a. Click the object icon .
- b. Select **Crypto Configuration > SSL Client Profile**
- c. Click **New** to open the SSL Client Profile window.

While you are configuring RBM:

- a. In the **Authenticate** section (if you have selected the LDAP method) or **Credential Mapping** section (if you have selected **Search LDAP for group name**) select an **SSL client type** of **Client profile**.
 - b. Click the plus icon  next to the **SSL client profile** field to open the SSL Client Profile window.
2. In the SSL Client Profile window, enter a name for your profile.
 3. Ensure that **Enable administrative state** is enabled, and optionally enter comments.
 4. Select which SSL and TLS protocols your profile supports.
 5. Specify which cipher suites your profile supports.
 6. Select from the following options:

Use SNI

Allows the client to send the Server Name Indication (SNI) extension in the ClientHello message to the server that the client attempts to connect to.

Permit connections to insecure SSL servers

Allows connections to SSL servers that do not support RFC 5746.

Enable compression

Enables SSL compression. Compression in HTTPS is dangerous because the HTTPS connection becomes vulnerable to the CRIME (Compression Ratio Info-leak Made Easy) attack.

7. Optionally specify that client connections will pass a host name in the SNI extension to the ClientHello, and specify the host name to pass.
8. Specify the client credentials. There are two parts to this:
 - Identification credentials specify the credentials that the appliance uses to authenticate itself to an SSL server if the SSL server requests client authentication.
 - Validation credentials are required if you select **Validate server certificate** and specify details about how the client authenticates the SSL server.

You can create the credentials profiles before you create the SSL client profile, and select them in the Identification credentials and Validation credentials list. Alternatively, you can click the plus icon  to open dialogs to create the two credentials.

9. Specify session caching features. Caching is enabled by default, and you can specify how long sessions are cached for in seconds and the minimum size of the cache in number of entries.
10. Optionally, open the **Advanced** section and add to the list of elliptic curves that the SSL client profile supports.

Related concepts

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Creating an SSL client profile by using the command line

Use the command line to create an SSL client profile.

Before you begin



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management” on page 440](#) before you attempt to set up authentication with LDAP

About this task

If you want to configure a secure SSL (TLS) connection with an LDAP server, then you must configure a client profile.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type `crypto` to enter crypto configuration mode.
4. Type the following command to configure an SSL client profile:

```
ssl-client name
```

Where *name* is a name for the profile you are creating.

5. Select which protocols the client profile supports:

```
protocols option-string
```

Where *option-string* Specifies the SSL and TLS protocol versions to support. When you enable support for multiple protocol versions, use a plus sign (+) character to separate the versions. The following values are valid. The default value is TLSv1d0+TLSv1d1+TLSv1d2.

- SSLv3 - Enables SSL 3.0.
- TLSv1d0 - Enables TLS 1.0.
- TLSv1d1 - Enables TLS 1.1.
- TLSv1d2 - Enables TLS 1.2.

6. Specify which cipher suites your client profile supports:

```
ciphers cipher-string
```

Where *cipher-string* specifies the supported cipher. You must repeat this command for each cipher that you support. See [“ciphers”](#) on page 961 for a list of available ciphers.

7. Optionally specify which elliptic curves your client supports:

```
curves name
```

Where *name* specifies the supported curve. You must repeat this command for each cipher that you support. See [“curves”](#) on page 962 for a list of available curves.

8. Specify options for your client profile:

```
ssl-client-features feature
```

Separate multiple features with the plus sign (+) character. The default value is `use-sni`. The following features are available:

use-sni

Allows the client to send the Server Name Indication (SNI) extension in the ClientHello message to the server that the client attempts to connect to.

permit-insecure-servers

Allows connections to SSL servers that do not support RFC 5746.

compression

Enables SSL compression. Compression in HTTPS is dangerous because the HTTPS connection becomes vulnerable to the CRIME (Compression Ratio Info-leak Made Easy) attack.

9. Specify the identification credentials for your client. These credentials are supplied to requesting SSL servers.

```
idcred name
```

Where *name* is the name of an existing identification credentials set (see [“idcred”](#) on page 733 for how to create an identification credentials set)

10. Optionally specify that the client will validate the SSL server, and specify the name of the validation credentials to use:

```
validate-server-cert on
valcred name
```

Where *name* is the name of an existing validation credentials set (see [“Validation Credentials commands”](#) on page 980 for how to create a validation credentials set).

11. Optionally specify that session caching is enabled, and specify the cache timeout and size:

```
caching on
cache-timeout seconds
cache-size entries
```

Example

The following example configures an SSL client profile named "myclient". The profile uses the default protocols and ciphers, and the default feature of enabling SNI. It uses the identity credentials set named "myclientcred", and enables server certificate validation by using the validation credentials set named "mcserver". The client caches a minimum of 150 session entries for up to ten minutes.

```
mqa# config
Global configuration mode
mqa(config)# crypto
Crypto configuration mode

mqa(config-crypto)# ssl-client myclient
mqa(config-crypto)# idcred myclientcred
mqa(config-crypto)# validate-server-cert on
mqa(config-crypto)# valcred mcserver
mqa(config-crypto)# caching on
mqa(config-crypto)# cache-timeout 600
mqa(config-crypto)# cache-size 150
mqa(config-crypto)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

Related concepts

[“Crypto commands”](#) on page 726

You can use the crypto commands to manage certificates on the IBM MQ Appliance.

[“SSL Client Profile commands”](#) on page 959

SSL Client Profile mode provides the commands to create or modify an SSL client profile.

[“Important: avoiding user lock out when configuring role based management”](#) on page 440

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Creating a load balancer group profile by using the web UI

You can create a profile describing a load balancer group of LDAP servers.

Before you begin



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management”](#) on page 440 before you attempt to set up authentication with LDAP

About this task

You can configure a load balancer group profile by using the IBM MQ Appliance web UI. You can do this by opening a dialog while you configure role based management (RBM).

Procedure

1. Open the Load Balancer Group profile window. In the **Authenticate** section (if you have selected the LDAP method) or **Credential Mapping** section (if you have selected **Search LDAP for group name**).

Click the plus icon  next to the **Load balancer group** field.

2. Enter a name for your load balancer group profile.
3. Ensure that **Enable administrative state** is selected and optionally enter some comments.
4. Select the algorithm for choosing between the LDAP servers in your load balancer group. The following options are available:

First Alive

Uses the concept of a primary server and 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.

Hash

Uses the IP address of the client as the basis for server selection.

With the hash algorithm, the same client is served by the same server. Use this algorithm only when clients access applications that require the storage of server-side state information, such as cookies. Hashing algorithms cannot ensure even distribution.

Least Connections

Maintains a record of active server connections and forward a new connection to the server with the least number of active connections.

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.

Weighted Round Robin

Maintains a weighted list of servers and forwards new connections in proportion to the weight (or preference) of each server.

5. Specify a damp time. The damp time is the period that a server is removed from the load balancer group because it cannot connect during a normal HTTP or TCP transaction. Such a server has a health state of `softdown`. When this interval expires, the server is restored to the load balancer group and placed in the up state. This command does not affect servers that are in the down state. Enter a value in the range 1 - 86400. The default value is 120.
6. Optionally select **Do not bypass down state** to block every connection when all members of the group are in the down state.
7. Optionally select **Try every server before failing**. When this option is selected, the appliance sends the request to each server until one responds or all fail. This command applies only when none of the group members are in the up state. Each server that fails is set to the `softdown` state.
8. Optionally select **Masquerade as group name** to pass the name of the load balancer group as the host name to the remote server.
9. Open the Health section and specify whether health checking is enabled. A health check is a scheduled rule that sends the same request to each member. The successful completion of the health check requires that the server passes normal TCP and HTTP connection criteria, depending on check type. See [“health-check” on page 851](#) for more information on the available options.
10. Open the Members section to specify details of the members of the load balancer group.
11. Click **Add** to add a member and supply the following information:

Actual host

The name or IP address of the server.

Weight

For weighted algorithms, specifies the relative weight (preference). Enter a value in the range 1 - 65000. The default value is 1.

Mapped server ports

Specifies the port on the real server. If nonzero, the associated real server is contacted on this port. Normally the real server is contacted on the same port number as the one for the virtual server. In this case, retain the default value of 0. If services run on different ports for different members of the group, define this value.

Health port

Specifies the port to test. Retain the default value of 0 to use the port that is defined for this load balancer group.

12. Click **Apply** to save your load balancer group profile.

Related concepts

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Creating a load balancer group profile by using the command line

Use the command line to create a load balancer group profile.

Before you begin



CAUTION: Read the topic [“Important: avoiding user lock out when configuring role based management” on page 440](#) before you attempt to set up authentication with LDAP

About this task

If you want to configure your LDAP user authentication or credential mapping to use a group of LDAP servers, you must configure a load balancer group profile on the appliance.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type `loadbalancer-group name` to create a load balancer group profile.
4. Select the algorithm that is used to select which member of the group to connect to:

```
algorithm selected-algorithm
```

Where *selected-algorithm* is one of the following options:

first-alive

Uses the concept of a primary server and 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.

hash

Uses the IP address of the client as the basis for server selection.

With the hash algorithm, the same client is served by the same server. Use this algorithm only when clients access applications that require the storage of server-side state information, such as cookies. Hashing algorithms cannot ensure even distribution.

least-connections

Maintains a record of active server connections and forward a new connection to the server with the least number of active connections.

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.

weighted-round-robin

Maintains a weighted list of servers and forwards new connections in proportion to the weight (or preference) of each server.

5. Specify the damp time. The damp time is the period that a server is removed from the load balancer group because it cannot connect during a normal HTTP or TCP transaction. Such a server has a health state of `softdown`. When this interval expires, the server is restored to the load balancer group and placed in the up state. This command does not affect servers that are in the down state.

```
damp seconds
```

Where *seconds* is the number of seconds that a server remains in a `softdown` state. Enter a value in the range 1 - 86400. The default value is 120.

6. Specify the action that the appliance takes when no member of the group is in the up state.

```
giveup-when-all-members-down
```

Specify a *setting* of on to not forward the connection to any member. The appliance makes the next connection attempt when at least one member is in the up state. Specify `off` to select the first member in the down state and forward the connection to this server. The default setting is `off`.

7. Optionally specify that the appliance sends the request to each server until one responds or all fail. This command applies only when none of the group members are in the up state. Each server that fails is set to the `softdown` state.

```
try-every-server setting
```

Where *setting* is on or off.

8. Optionally specify that the name of the load balancer group is passed as the host name to the remote server.

```
masquerade setting
```

Where *setting* is on or off.

9. Specify whether a health check is to be implemented on members of the group, and configure it.

```
health-check options
```

For the options available and more information, see [“health-check” on page 851](#)

10. Specify the LDAP servers that belong to the load balancer group:

```
server address [weight] [mapped-port] [health-port]
```

address

The name or IP address of the server.

weight

For weighted algorithms, specifies the relative weight (preference). Enter a value in the range 1 - 65000. The default value is 1.

mapped-port

Specifies the port on the real server. If nonzero, the associated real server is contacted on this port. Normally the real server is contacted on the same port number as the one for the virtual server. In this case, retain the default value of 0. If services run on different ports for different members of the group, define this value.

health-port

Specifies the port to test. Retain the default value of 0 to use the port that is defined for this load balancer group.

Example

The following example configures a load balancer group profile named "LBGroup". The profile specifies a first-alive algorithm, takes the defaults for other settings, and specifies the three LDAP servers that comprise the load balancing group.

```
mqa# config
Global configuration mode
mqa(config)# loadbalancer-group LBGroup
mqa(config loadbalancer-group LBgroup)# algorithm first-alive
mqa(config loadbalancer-group LBgroup)# server ldap1.here.com
mqa(config loadbalancer-group LBgroup)# server ldap2.here.com
mqa(config loadbalancer-group LBgroup)# server ldap3.here.com
mqa(config loadbalancer-group LBgroup)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

Related concepts

[“Load Balancer Group commands” on page 849](#)

Load Balancer Group mode provides the commands to create or modify a load balancer group configuration.

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

User authentication with XML file

You can configure the IBM MQ Appliance to authenticate users by using an XML file.

You create an XML file that contains user authentication details. You can use an XML file builder that is provided on the appliance to create the file.

After you have configured how users are authenticated by using an XML file, you must go on to specify how authenticated users are authorized to use the appliance resources. You do this by configuring credential mapping. You can use the same XML file for both authentication and credential mapping. Alternatively, you can perform credential mapping by using credentials that are defined for the group that the user belongs to.

You can use the RBM builder in the IBM MQ Appliance web UI to create the XML file. If you want to create the file manually, an example file, `store:///RBMInfo.xml`, is provided to guide you. The final RBM file must conform to the `store:///AAAIInfo.xsd` schema.

Related concepts

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Configuring user authentication with an XML file by using the web UI

Configure the appliance to authenticate the users defined in an XML file.

About this task

You can use the IBM MQ Appliance web UI to configure role based management such that the appliance uses user credentials that are defined in an XML file.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > RBM Settings**.
3. Ensure that **Enable administrative state** is selected (it is selected by default) and click **Authentication** to view the authentication options.
4. Select an **Authentication method** of **XML file**.
5. If you already defined an XML file, enter its path in **XML file URL**. Alternatively you can edit an existing file, or create a new file, by using the RBM builder. Click **Edit** or **New**. See [“Using the RBM builder to create an XML file” on page 458](#) for help with using the RBM builder.
6. To define fallback users, you can choose **All users** from the **Local accounts for fallback** list to have all local users able to log in to the appliance if the XML file is unavailable. Alternatively, select **Specific users** and select one or more local users.
7. Optionally, change the default cache settings. Cache settings determine how long user details are held on the appliance before authentication is performed again. By default, the appliance retains details for an absolute period of 600 seconds. You can change the cache mode or the cache lifetime, or both. You can also disable caching altogether.
8. Click **Apply** to apply your changes.

What to do next

After you specify the user authentication method, you must next configure credential mapping.

Related concepts

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Configuring user authentication with an XML file by using the command line

Use the command line to configure the appliance to authenticate the users defined in an XML file.

About this task

You can use commands to configure role based management such that the appliance uses user credentials that are defined in an XML file. You must have already created the XML file (see [“User authentication with XML file” on page 455](#))

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `confi`g to enter global configuration mode.
3. Type the following command to configure role based management:

```
ibm
```

4. Enter the following command to specify the XML file authentication method:

```
au-method xmlfile
```

5. Specify the URL of the XML you want to use.

```
au-info-url URL
```

6. Optionally specify fallback users who can log in to the appliance if the XML file is not available. Fallback users must already have been added as local users to the appliance. You can specify that all local users are fallback users by entering the following command:

```
fallback-login local
```

Alternatively, you can specify one or more particular users by entering the following commands:

```
fallback-login restricted
fallback-user localuser1
fallback-user localuser2
...
fallback-user localuserN
```

7. Alter the default cache settings, if required. By default, the appliance caches results of authentication attempts for 600 seconds, but you can change the mode of caching, and the caching duration by entering the following commands:

```
au-cache-mode mode
au-cache-ttl seconds
```

Where mode is one of:

absolute

Caches the results of user authentications for a period of time specified by the **au-cache-ttl** command (the explicit time-to-live). This is the default setting.

disabled

Disables caching. The appliance will not cache any results and instead always authenticates every time a user requests access.

maximum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the maximum of the two values.

minimum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the minimum of the two values.

Example

The following example configures the appliance to use the authentication details defined in the file `store:///RBMinfo.xml`.

```
mqa# config
Global configuration mode
mqa(config)# rbm
Modify RBM Settings configuration

mqa(config rbm)# au-method xmlfile
mqa(config rbm)# au-info-url store:///RBMinfo.xml
mqa(config rbm)# fallback-login local
mqa(config rbm)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

What to do next

After you specify the user authentication method, you must next configure credential mapping.

Related concepts

[“RBM setting commands” on page 917](#)

Use the RBM commands to configure role based management on the IBM MQ Appliance.

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

Using the RBM builder to create an XML file

You can use the RBM builder in the IBM MQ Appliance web UI to define an XML file. You can use the file for user authentication or credential mapping, or for both.

About this task

To use the RBM builder:

- If you are using the XML file for user authentication:
 1. In the Authentication section of the RBM settings, select an **Authentication method of XML file**.
 2. Click **New** for the **XML file URL** to open the builder.
- If you are using the XML file for user authorization:
 1. In the Credential-mapping section of the RBM settings, select a **Credential-mapping method of XML file**.
 2. Click **New** for the **XML file URL** to open the builder.
- If you are using the XML for both user authentication and user authorization:
 1. In the Authentication section of the RBM settings, select an **Authentication method of XML file**.
 2. Click **New** for the **XML file URL** to open the builder.
 3. After you have created the XML file, in the Credential-mapping section of the RBM settings, select a **Credential-mapping method of XML file**, and specify the URL of the file that you created.

Procedure

To create an XML file by using the RBM builder:

1. In the first Edit RBM Policy file page, click **Next**.
2. In the Default Credentials page, enter the name of a default credentials profile, if required. Any user that fails authentication is granted this credential. Leave blank to deny access to all users who fail authentication. Click **Next**.
3. In the User Identities page, click **Next** if you are using the file for user authorization only. Otherwise, click **Add**. Enter a user name, user password, and the name of the credential profile that you want to assign to this user. Click **Submit**. Repeat for every user that you want to add, then click **Next**.
4. In the Access Profile Mappings page, click **Next** if you are using the file for user authentication only. Otherwise, click **Add**. Enter a name for the profile you are creating, then click **Build** to specify the permissions for the policy. Click **Help** in the Builder window for guidance on defining the policy.
5. In the next page, specify the name of the XML file you are creating, and the local directory to store it in. Click **Next**.
6. In the Confirm Creation page, click **Commit**.

Related concepts

[“Important: avoiding user lock out when configuring role based management” on page 440](#)

You must take care when you configure role based management (RBM) that you do not make it possible for all users to be locked out of the appliance.

User authentication with local users

You can configure the IBM MQ Appliance to authenticate users who are locally configured on the appliance.

Appliance users are created by using the appliance user commands or the Appliance section of the IBM MQ Appliance web UI

Local users can belong to locally defined user groups. User privileges are then defined according to the group that they belong to.

Configuring user authentication with local users by using the web UI

Configure the appliance to authenticate the local users defined on the appliance.

About this task

You can use the IBM MQ Appliance web UI to configure role based management such that the appliance uses local definitions. You define the local users in a separate procedure, see [“Configuring user authentication with local users by using the web UI”](#) on page 459.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > RBM Settings**
3. Ensure that **Enable administrative state** is selected (it is selected by default) and click **Authentication** to view the authentication options.
4. Select an **Authentication method** of **Local user**.
5. Optionally, change the default cache settings. Cache settings determine how long user details are valid before authentication is performed again. By default, the appliance retains details for an absolute period of 600 seconds. You can change the cache mode or the cache lifetime, or both. You can also disable caching altogether.
6. Click **Apply** to apply your changes.

What to do next

After you specify the user authentication method, you must next configure credential mapping.

Configuring user authentication with local users by using the command line

Use the command line to configure the appliance to authenticate local users defined on the appliance.

About this task

You can use commands to configure role based management such that the appliance uses local user definitions. You create the users in a separate procedure, see [“Configuring local users by using the command line”](#) on page 472.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure role based management:

```
rbm
```

4. Enter the following command to specify the XML file authentication method:

```
au-method local
```

5. Alter the default cache settings, if required. By default, the appliance caches results of authentication attempts for 600 seconds, but you can change the mode of caching, and the caching duration by entering the following commands:

```
au-cache-mode mode  
au-cache-ttl seconds
```

Where mode is one of:

absolute

Caches the results of user authentications for a period of time that is specified by the **au-cache-ttl** command (the explicit time-to-live). This setting is the default setting.

disabled

Disables caching. The appliance will not cache any results and instead always authenticates every time a user requests access.

maximum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the maximum of the two values.

minimum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the minimum of the two values.

Example

The following example configures the appliance to use local user definitions.

```
mqa# config  
Global configuration mode  
mqa(config)# rbm  
Modify RBM Settings configuration  
  
mqa(config rbm)# au-method local  
mqa(config rbm)# exit  
mqa(config)# write memory  
Overwrite previously saved configuration? Yes/No [y/n]: y  
Configuration saved successfully.
```

What to do next

After you specify the user authentication method, you must next configure credential mapping.

Related concepts

[“RBM setting commands” on page 917](#)

Use the RBM commands to configure role based management on the IBM MQ Appliance.

Credential mapping with an XML file

You can configure the IBM MQ Appliance to map user names onto access policies that are defined in an XML file.

You create an XML file that contains policy details. You can use an XML file builder that is provided on the appliance to create the file. The same file can contain user authentication details.

The policy defines the set of permissions that a user has to the various appliance resources. This feature is known as user authorization. A user name is mapped onto these policies according to the user authentication method:

- If users are authenticated by using LDAP, then the user's distinguished name is the input credential for policies that are defined in the XML file. If the appliance is configured to look up groups the user

belongs to in an LDAP server, then the returned attribute that is specified in the LDAP search parameters is used as the input credential for each group instead.

- If users are authenticated by using an XML file, then the same user credential is used as input for policies that are defined in the XML file.
- If you use local user definitions, then the user name is the input credential for policies that are defined in the XML file.

You can use the RBM builder in the IBM MQ Appliance web UI to create the XML file. If you want to create the file manually, an example file, `store:///RBMInfo.xml`, is provided to guide you. The final RBM file must conform to the `store:///AAAIInfo.xsd` schema.

You can, if required, configure an LDAP search to retrieve user groups from LDAP directories for XML file or local authenticated users. Returned user groups can then be mapped onto access policies defined in the XML file. In these cases you must configure the LDAP search so that the XML file or local user name is used as the input credential to the LDAP query. You might then need, for example, to append a common suffix to the user name to build an LDAP user distinguished name when querying the user's group membership.

Configuring credential mapping with an XML file by using the web UI

Configure the appliance to authorize users by using access policies that are defined in an XML file.

About this task

You can use the IBM MQ Appliance web UI to configure role based management such that the appliance uses access policies that are defined in an XML file.

You can specify that an LDAP directory is searched for groups that the authenticated user belongs to, then the returned groups are mapped onto access policies in the XML file. The LDAP query should search for groups that the user belongs to. Do not configure a search that looks for users in a particular group; if the search returns users you will be attempting to map groups onto users, rather than users onto groups.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > RBM Settings**
3. Ensure that **Enable administrative state** is selected (it is selected by default) and click **Credential-mapping** to view the authentication options.
4. Select an **Credential-mapping method** of **XML file**.
5. If you already defined an XML file, enter its path in **XML file URL**. Alternatively you can edit an existing file, or create a new file, by using the RBM builder. Click **Edit** or **New**. See [“Using the RBM builder to create an XML file” on page 458](#) for help with using the RBM builder.
6. If you have defined a user authentication method of LDAP, select **Search LDAP for group name** to look up an attribute (usually the distinguished name) of each group the user belongs to. These attributes are then used as the input credential to the XML file. Otherwise, the distinguished name of the user is used as the input credential. If you select this option, you must then supply the details for connecting to the LDAP server:
 - a) Specify the **Server host** and the **Server port** for connecting to the LDAP server (server port is usually 389, or 636 for an SSL connection).
 - b) If you have configured a load balancer group for LDAP access and created a profile, specify it in the **Load balancer group** field. Alternatively, click the plus icon  to open the Load Balancer Group dialog to specify a profile for your load balancer group, see [“Creating a load balancer group profile by using the web UI” on page 451](#). You can leave this field blank if you are using a single LDAP server.
 - c) Specify the DN that the appliance uses to bind to the LDAP server to perform the search in the **LDAP bind DN** field. Specify the password alias in the **LDAP bind password alias** field. Click the

plus icon  to create a password alias if you have not already created one. (Leave these fields blank if you are using an anonymous bind to access the LDAP server.)

- d) Specify the **LDAP search parameters**. You can enter these parameters directly, or you can click the plus icon to open the LDAP Search Parameters dialog. Your search must look for the user group or groups that the authenticated user belongs to, and return one or more user group names.
- e) Specify an **LDAP read timeout**. The timeout is the time that the appliance will attempt to connect to the LDAP server before closing the connection. The default is 60 seconds. Specify 0 to never timeout.
- f) If you want to use an SSL (TLS) connection to the LDAP server, select an **SSL client type** of **Client profile**. If you have already defined a profile, select the profile name from the **SSL client profile** list. Alternatively, click the plus icon  to open the SSL Client Profile dialog and create a new SSL client profile, see [“Creating an SSL client profile by using the web UI” on page 448](#).

7. Click **Apply** to apply your changes.

What to do next

After you specify the credential mapping method, you can next configure the password policy for your local users (password policy does not apply to other user types).

Configuring credential mapping with an XML file by using the command line

Use the command line to configure the appliance to authorize users by using access policies that are defined in an XML file.

About this task

You can use commands to configure role based management such that the appliance uses access policies that are defined in an XML file. You must have already created the XML file (see [“Credential mapping with an XML file” on page 460](#))

You can specify that an LDAP directory is searched for groups that the authenticated user belongs to, then the returned groups are mapped onto access policies in the XML file. The LDAP query should search for groups that the user belongs to. Do not configure a search that looks for users in a particular group; if the search returns users you will be attempting to map groups onto users, rather than users onto groups.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure role based management:

```
rbm
```

4. Enter the following command to specify the XML file authorization method:

```
mc-method local
```

5. Specify the URL of the XML file that you want to use.

```
mc-info-url URL
```

6. If you have defined a user authentication method of LDAP, you can look up an attribute (usually the distinguished name) of each group the user belongs to. These attributes are then used as the input credential to the XML file. Otherwise, the distinguished name of the user is used as the input credential. If you want to perform an LDAP search, you must then supply the details for connecting to the LDAP server:
 - a) Specify that you want to perform an LDAP search:

```
mc-ldap-search on
```

- b) Specify the server and password for connecting to the LDAP server:

```
mc-server-host host  
mc-server-port port
```

Where *host* is the IP address or domain name of the LDAP server and *port* is usually 389, or 636 for an SSL connection.

- c) If you have configured a load balancer group for LDAP access and created a profile, specify the profile by using the following command:

```
mc-loadbalancer-group name
```

Where *name* is the name of the load balancer group profile.

- d) Specify the Distinguished Name and password that the appliance uses to bind to the LDAP server to perform the search by using the following commands. (Omit these commands if you are using an anonymous bind to access the LDAP server.)

```
mc-ldap-bind-dn dn  
mc-ldap-bind-password password
```

Where *dn* is the Distinguished Name and *password* is the password.

- e) Specify the LDAP search parameters by entering the following command:

```
mc-ldap-parameters parameters_name
```

Where *parameters_name* is the name of a set of LDAP search parameters that you have previously defined using the **ldap-search-parameters** command. Your search must look for the user group or groups that the authenticated user belongs to, and return one or more user group names.

- f) Specify a timeout for the LDAP search. This is the time that the appliance will attempt to connect to the LDAP server before closing the connection. The default is 60 seconds. Specify 0 to never timeout.

```
mc-ldap-readtimeout timeout
```

Where *timeout* specifies the duration of the timeout period.

- g) If you want to use an SSL (TLS) connection to the LDAP server, specify an SSL client profile to use by entering the following commands:

```
ssl-client-type client  
ssl-client client_profile
```

Where *client_profile* is the name of a client profile that you have previously created.

Example

The following example configures the appliance to use the authorization details defined in the file `store:///RBMinfo.xml`.

```
mqa# config  
Global configuration mode  
mqa(config)# rbm  
Modify RBM Settings configuration  
  
mqa(config rbm)# mc-method xmlfile  
mqa(config rbm)# mc-info-url store:///RBMinfo.xml  
mqa(config rbm)# exit  
mqa(config)# write memory  
Overwrite previously saved configuration? Yes/No [y/n]: y  
Configuration saved successfully.
```

The following commands configure the appliance to use the authorization details defined in the file `store:///RBInfo.xml`, performing an LDAP search to retrieve user group names from an LDAP repository:

```
mqa# config
Global configuration mode
mqa(config)# rbm
Modify RBM Settings configuration

mqa(config rbm)# mc-method xmlfile
mqa(config rbm)# mc-info-url store:///RBInfo.xml
mqa(config rbm)# mc-ldap-search on
mqa(config rbm)# mc-server-host LDAP_serv1
mqa(config rbm)# mc-server-port 389
mqa(config rbm)# mc-ldap-bind-dn "cn=proxyuser"
mqa(config rbm)# mc-ldap-bind-password p@Ssw0rd
mqa(config rbm)# mc-ldap-parameters ldap1-MC
mqa(config rbm)# mc-ldap-readtimeout 120
mqa(config rbm)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

What to do next

After you specify the credential mapping method, you can next configure the password policy for your local users (password policy does not apply to other user types).

Related concepts

“RBM setting commands” on page 917

Use the RBM commands to configure role based management on the IBM MQ Appliance.

Credential mapping with local user groups

You can configure the IBM MQ Appliance to use local user groups for user authorization.

You create user groups locally on the appliance, and define what access members of that group have to appliance resources. This access is known as user authorization. A user name is mapped onto a user group according to the user authentication method:

- If users are authenticated by using LDAP, then the user's distinguished name is used to perform a further LDAP query (in the same directory or a different one). The second query retrieves the user group or groups for that distinguished name, and these groups are in turn mapped onto local user groups. If, as a result of this process, multiple local user groups apply to the user, then the access profiles for the groups are combined so the user has the superset of the authority they grant. If a local group is not found for a particular LDAP group then the group is ignored. If no LDAP groups are returned for the user, or no matching local user groups are found, then the user has no authority and access is denied.
- If users are authenticated by using an XML file, then the user credential profiles specified for the user in the file is mapped onto local user groups.
- If you use local user definitions, then the local groups defined for that user are used for authorization.

You can, if required, configure an LDAP search to retrieve user groups from LDAP directories for XML file or local authenticated users. Returned user groups can then be mapped onto local user groups. In these cases you must configure the LDAP search so that the XML file or local user name is used as the input credential to the LDAP query. You might then need, for example, to append a common suffix to the user name to build an LDAP user distinguished name when querying the user's group membership.

Configuring credential mapping with local user groups by using the Web UI

Configure the appliance to authorize users by using local user group definitions.

About this task

You can use the IBM MQ Appliance web UI to configure role based management such that the appliance uses local user groups for user authorization. You define the local users in a separate procedure, see [“Configuring local user groups by using the web UI” on page 474](#).

You can specify that an LDAP directory is searched for groups that the authenticated user belongs to, then the returned groups are mapped onto local user groups. The LDAP query should search for groups that the user belongs to. Do not configure a search that looks for users in a particular group; if the search returns users you will be attempting to map groups onto users, rather than users onto groups.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > RBM Settings**
3. Ensure that **Enable administrative state** is selected (it is selected by default) and click **Credential-mapping** to view the authentication options.
4. Select an **Credential-mapping method** of **Local user group**.
5. If you have defined a user authentication method of LDAP, then you must select **Search LDAP for group name**. You must then supply the details for connecting to the LDAP server:
 - a) Specify the **Server host** and the **Server port** for connecting to the LDAP server (server port is usually 389, or 636 for an SSL connection).
 - b) If you have configured a load balancer group for LDAP access and created a profile, specify it in the **Load balancer group** field. Alternatively, click the plus icon  to open the Load Balancer Group dialog to specify a profile for your load balancer group.
 - c) Specify the DN that the appliance uses to bind to the LDAP server to perform the search in the **LDAP bind DN** field. Specify the password alias in the **LDAP bind password alias** field. Click the plus icon  to create a password alias if you have not already created one. (Leave these fields blank if you are using an anonymous bind to access the LDAP server.)
 - d) Specify the **LDAP search parameters**. You can enter these directly, or you can click the plus icon to open the LDAP Search Parameters dialog. Your search must look for the user group or groups that the authenticated user belongs to, and return one or more user group names.
 - e) Specify an **LDAP read timeout**. This is the time that the appliance will attempt to connect to the LDAP server before closing the connection. The default is 60 seconds. Specify 0 to never timeout.
 - f) If you want to use an SSL (TLS) connection to the LDAP server, select an **SSL client type** of **Client profile**. If you have already defined a profile, select the profile name from the **SSL client profile** list. Alternatively, click the plus icon  to open the SSL Client Profile dialog and create a new SSL client profile.
6. Click **Apply** to apply your changes.

What to do next

After specifying the credential mapping method, you can next configure the password policy for your local users (password policy does not apply to other user types).

Configuring credential mapping with local user groups by using the command line

Use the command line to configure the appliance to authorize users by using local user group definitions.

About this task

You can use commands to configure role based management such that the appliance uses local user groups for user authorization. You must create user groups as a separate operation.

You can specify that an LDAP directory is searched for groups that the authenticated user belongs to, then the returned groups are mapped onto local user groups. The LDAP query should search for groups that the user belongs to. Do not configure a search that looks for users in a particular group; if the search returns users you will be attempting to map groups onto users, rather than users onto groups.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure role based management:

```
rbm
```

4. Enter the following command to specify the local user group authorization method:

```
mc-method local
```

5. If you are using local user groups to authorize LDAP users, define an LDAP search to retrieve the user group for the authenticated user, that can in turn be mapped onto a local user group. Enter the following commands:

- a) Specify that you want to perform an LDAP search:

```
mc-ldap-search on
```

- b) Specify the server and password for connecting to the LDAP server:

```
mc-server-host host  
mc-server-port port
```

Where *host* is the IP address or domain name of the LDAP server and *port* is usually 389, or 636 for an SSL connection.

- c) If you have configured a load balancer group for LDAP access and created a profile, specify the profile by using the following command:

```
mc-loadbalancer-group name
```

Where *name* is the name of the load balancer group profile.

- d) Specify the Distinguished Name and password that the appliance uses to bind to the LDAP server to perform the search by using the following commands. (Omit these commands if you are using an anonymous bind to access the LDAP server.)

```
mc-ldap-bind-dn dn  
mc-ldap-bind-password password
```

Where *dn* is the Distinguished Name and *password* is the password.

- e) Specify the LDAP search parameters by entering the following command:

```
mc-ldap-parameters parameters_name
```

Where *parameters_name* is the name of a set of LDAP search parameters that you have previously defined using the **ldap-search-parameters** command. Your search must look for the user group or groups that the authenticated user belongs to, and return one or more user group names.

- f) Specify a timeout for the LDAP search. This is the time that the appliance will attempt to connect to the LDAP server before closing the connection. The default is 60 seconds. Specify 0 to never timeout.

```
mc-ldap-readtimeout timeout
```

Where *timeout* specifies the duration of the timeout period.

- g) If you want to use an SSL (TLS) connection to the LDAP server, specify an SSL client profile to use by entering the following commands:

```
ssl-client-type client  
ssl-client client_profile
```

Where *client_profile* is the name of a client profile that you have previously created.

6. Save your configuration and exit.

Example

The following example configures the appliance to use the authorization details defined by local user groups.

```
mqa# config
Global configuration mode
mqa(config)# rbm
Modify RBM Settings configuration

mqa(config rbm)# mc-method local
mqa(config rbm)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

The following commands configure the appliance to use the authorization details defined by local user groups, performing an LDAP search to retrieve user group names from an LDAP repository:

```
mqa# config
Global configuration mode
mqa(config)# rbm
Modify RBM Settings configuration

mqa(config rbm)# mc-method local
mqa(config rbm)# mc-ldap-search on
mqa(config rbm)# mc-server-host LDAP_serv1
mqa(config rbm)# mc-server-port 389
mqa(config rbm)# mc-ldap-bind-dn "cn=proxyuser"
mqa(config rbm)# mc-ldap-bind-password p@Ssw0rd
mqa(config rbm)# mc-ldap-parameters ldap1-MC
mqa(config rbm)# mc-ldap-readtimeout 120
mqa(config rbm)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

What to do next

After specifying the credential mapping method, you can next configure the password policy for your local users (password policy does not apply to other user types).

Related concepts

[“RBM setting commands” on page 917](#)

Use the RBM commands to configure role based management on the IBM MQ Appliance.

Password policy

You can define a password policy as part of your RBM configuration.

The password policy applies only to locally defined users. It does not apply to users who are defined externally or in an XML file.

The password policy enables you to dictate the requirements for a password, such as minimum length, whether it must contain some numbers, some non-alphanumeric characters, some mixed-case characters and so on. You can also specify how often the password must be changed.

Configuring a password policy by using the web UI

Configure the password policy for locally defined users.

About this task

You can use the IBM MQ Appliance web UI to configure the password policy for users defined locally on the appliance. The policy does not apply to users defined in other ways.

Procedure



1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > RBM Settings**
3. Ensure that **Enable administrative state** is selected (it is selected by default) and click **Password policy**.
4. Define the policy by specifying the following options:
 - a) In the **Minimum length** field, enter the minimum number of characters in a password.
 - b) Define characteristics for passwords:
 - Set the **Require mixed case** property to require mixed case passwords.
 - Set the **Require non-alphanumeric** property to require nonalphanumeric characters in passwords.
 - Set the **Require number** property to require numeric characters in passwords.
 - Set the **Disallow user name substring** property to indicate whether to allow the user name string in the password. If the user is `george`, the property controls whether to allow `george1!` or `passgeorgeword` as the password.
 - c) Set the **Enable aging** property to control password-aging. If enabled, define the maximum password age in the **Maximum age** field.
 - d) Set the **Control reuse** property to control the reuse of previous passwords. If enabled, define the reuse history by entering the number of past passwords to compare against for reuse in the **Reuse history** field.
 - e) From the **Password hash** algorithm list, select the algorithm that is used to hash passwords before they are stored.
5. Click **Apply** to apply your changes.

What to do next

After you specify a password policy, you can next configure an account policy.

Configuring a password policy by using the command line

Use the command line to configure a password policy for local users.

About this task

You can use commands to configure the password policy for users defined locally on the appliance. The policy does not apply to users defined in other ways.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure role based management:

```
rbm
```

4. Use the following commands, as required, to configure your password policy:

- Enter the following command to specify the minimum length of the password:

```
pwd-minimum-length length
```

- Enter the following command to specify that passwords must contain both uppercase and lowercase characters:

```
pwd-mixed-case on
```

- Enter the following command to specify that the password must contain non-alphanumeric characters in addition to alphanumeric characters:

```
pwd-nonalpha numeric on
```

- Enter the following command to specify that the password must contain numbers:

```
pwd-digit on
```

- Enter the following command to indicate whether to allow the user name string in the password:

```
pwd-username on
```

- Enter the following commands to specify password aging in the policy:

```
pwd-aging on  
pwd-max-age days
```

Where *days* is the number of days before the password expires.

- Enter the following commands to specify that the reuse of recent passwords should be controlled:

```
pwd-history on  
pwd-max-history count
```

Where *count* is the number of passwords to retain.

- Enter the following command to specify the hash algorithm to use when encrypting the password:

```
password-hash-algorithm algorithm
```

Where *algorithm* is one of:

- md5crypt (default)
- sha256crypt

Example

The following example configures a password policy that requires user passwords are 8 characters long, must contain a number, and expires after three months.

```
mqa# config  
Global configuration mode  
mqa(config)# rbm  
Modify RBM Settings configuration  
  
mqa(config rbm)# pwd-minimum-length 8  
mqa(config rbm)# pwd-digit on  
mqa(config rbm)# pwd-aging on  
mqa(config rbm)# pwd-max-age 90  
mqa(config rbm)# exit  
mqa(config)# write memory  
Overwrite previously saved configuration? Yes/No [y/n]: y  
Configuration saved successfully.
```

What to do next

After you specify a password policy, you can next configure an account policy.

Related concepts

[“RBM setting commands” on page 917](#)

Use the RBM commands to configure role based management on the IBM MQ Appliance.

Account policy

You can define an account policy as part of your RBM configuration.

The account policy applies only to locally defined users. It does not apply to users who are defined externally or in an XML file.

The account policy enables you to specify how many failed log in attempts can occur before you lock out that connection, and how long the lock out remains in force. You can also specify how long a CLI session can be idle before it times out.

You can also restrict the `admin` account so it can access the appliance only by using the serial interface.

Configuring an account policy by using the web UI

Configure the account policy for locally defined users.

About this task

You can use the IBM MQ Appliance web UI to configure the account policy for users defined locally on the appliance. The policy does not apply to users defined in other ways.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > RBM Settings**
3. Ensure that **Enable administrative state** is selected (it is selected by default) and click **Account policy**.
4. Define the policy by specifying the following options:
 - a) Select **Restrict admin to serial** to specify that the `admin` account can connect only by using the serial interface.
 - b) Specify the **Maximum failed log ins** that a connection can have before it times out. The default is 0 to indicate that there is no restriction.
 - c) Specify the time in minutes that a lock out should last for. The default is 1 minute.
 - d) Specify the time that a CLI session can be idle before it times out. The default is 0 seconds to indicate that there is no timeout.
5. Click **Apply** to apply your changes.

Configuring an account policy by using the command line

Use the command line to configure an account policy for local users.

About this task

You can use commands to configure the account policy for users defined locally on the appliance. The policy does not apply to users defined in other ways.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure role based management:

```
ibm
```

4. Use the following commands, as required, to configure your account policy:

- Specify that you want to restrict the admin account to connect only by using the serial port by entering the following command:

```
restrict-admin on
```

- Specify the maximum number of login attempts that can occur before an account is locked out:

```
max-login-failure count
```

- Specify the duration to lock out accounts for after the specified number of failed logins:

```
lockout-duration minutes
```

- Enter the following command to specify a CLI timeout:

```
cli-timeout seconds
```

Example

The following example configures an account policy that specifies a user can have three attempts to log in before they are locked out for an hour. CLI sessions timeout if they are inactive for twenty minutes.

```
mqa# config
Global configuration mode
mqa(config)# rbm
Modify RBM Settings configuration

mqa(config rbm)# max-login-failure 3
mqa(config rbm)# lockout-duration 60
mqa(config rbm)# cli-timeout 1200
mqa(config rbm)# exit
mqa(config)# write memory
Overwrite previously saved configuration? Yes/No [y/n]: y
Configuration saved successfully.
```

Related concepts

[“RBM setting commands” on page 917](#)

Use the RBM commands to configure role based management on the IBM MQ Appliance.

Local users and user groups

You can define local users and local user groups on your appliance

You can configure role based management so that user authentication is set up for local users. You can also specify that local users are used as fall back when a remote authentication platform is unavailable (ensuring that someone can always log on to the appliance). You can separately specify that local user groups are used for credential mapping. You can use this setting for local user authentication, or any of the other authentication methods available.

Appliance users are created by using the appliance user commands or the Appliance section of the IBM MQ Appliance web UI

Local users can belong to locally defined user groups. User privileges are then defined according to the group that they belong to.

Configuring local users by using the web UI

You can configure local users by using the IBM MQ Appliance web UI.

About this task

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 the Administration icon .
2. Select **Access > User Account**.
3. Click **New** to add a user.
4. Specify the name of the user that you are configuring. This is the name that the local user uses to log in to the appliance. 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. Optionally select **Suppress initial password change**. Otherwise, the user is requested to change their password the first time that they log in.
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:
 - a) Click **Add**.
 - b) Specify the **Engine ID**, this ID is always 0 for the IBM MQ Appliance.
 - c) Select the authentication protocol that is used from the list. This is one of **SHA, MD5, or none**.
 - d) If you have selected a protocol other than **none**, specify whether authentication uses a key or a plain-text password.
 - e) Enter the password, 16-byte key (MD5 protocol), or 20-byte key (SHA protocol) in the **Authentication secret** field. Enter a key in hexadecimal notation.
 - f) Select the encryption type used from the **Privacy Protocol** list. The type is one of **AES, DES, or none**.
 - g) If you have selected a protocol other than **none**, specify whether encryption uses a key or a plain-text password.
 - h) Enter the password or key in the **Privacy secret** field. Enter a key in hexadecimal notation.
9. Click **Apply** to save the new user configuration.

Configuring local users by using the command line

You can configure local 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 207. Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure a local 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. Specify a password for the user:

```
password password
```

5. Set the user level of the account:

```
access-level level
```

Where *level* is `group-defined` to specify privileges according to group, or `privileged` to define a user with all privileges (that is, and administrative user).

6. Specify the group that the user belongs to, if required:

```
group groupName
```

7. Optionally, specify SNMPv3 credentials for the user:

```
snmp-cred engine-ID authentication-protocol authentication-secret-type authentication-secret  
privacy-protocol privacy-secret-type privacy-secret
```

Where:

engine-ID

This ID is set to 0 to identify the local engine.

authentication-protocol

Specify `none` to use no authentication key, `md5` to use HMAC-MD5-96 as the authentication protocol, or `sha` to use HMAC-SHA-96 as the authentication protocol.

authentication-secret-type

Enter `password` to specify that the authentication secret is a password that is converted to an intermediate key with a standardized algorithm. Enter `key` to specify a fully localized key. This parameter is not required if you have set `authentication-protocol` to `none`.

authentication-secret

Specifies the secret, or key, for authentication for this account.

- If a password, specify a plaintext password that is at least 8 characters long.
- If a key and HMAC-MD5 are the authentication protocol, specify the hex representation of a 16-byte key.
- If a key and HMAC-SHA-96 are the authentication protocol, specify the hex representation of a 20-byte key.

This parameter is not required if you have set *authentication-protocol* to none.

privacy-protocol

Identifies which privacy (encryption) protocol to use. Set to none to not encrypt data. Set to des to use CBC-DES (this setting is the default). Set to aes to use CFB128-AES-128.

privacy-secret-type

Enter password to specify that the privacy secret is a password that is converted to an intermediate key with a standardized algorithm. Enter key to specify a fully localized key.

privacy-secret

Specifies the secret, or key, for privacy (encryption) for this account. This parameter is required when the privacy protocol is des or aes.

- If a password, specify a plain text password that is at least 8 characters long.
- If a key, and HMAC-MD5 is the authentication protocol, specify the hex representation of a 16-byte key.
- If a key, and HMAC-SHA-96 is the authentication protocol, specify the hex representation of a 20-byte key.

8. After you configure the user, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Related reference

[“password” on page 721](#)

This command manages the password for an account.

[“access-level” on page 720](#)

This command assigns an account type of an account.

[“group” on page 720](#)

This command associates an account with a user group.

[“snmp-cred” on page 721](#)

This command adds SNMP V3 credentials to the account.

Configuring local user groups by using the web UI

You can configure local user groups by using the IBM MQ Appliance web UI.

About this task

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 the Administration icon .
2. Select **Access > 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. Specify an access profile for the group. You can click **Build** to open the access policy. You can add several policies, if required.
 7. Click **Apply** to save the new user configuration.

Configuring local user groups by using the command line

You can configure local user groups by using the command line interface.

About this task

To configure a local 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. You define the local users in a separate procedure, see [“Configuring local user groups by using the command line” on page 475](#)

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 207](#). Log in as an administrative user.
2. Type `config` to enter global configuration mode.
3. Type the following command to configure a local user group:

```
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. Set the access policy for the user group:

```
access-policy statement
```

Where *statement* defines the access policy that applies to users who belong to this group. A policy statement takes the 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

Enter `*` to match all domains.

resource

The resource type to which this policy applies. The special value * matches all resource types.

Name=name

Optionally 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.

5. After you configure the user group, enter `exit` to save the configuration and exit, or type `cancel` to exit without saving.

Routine user administration

Various routine administration tasks can be completed for local users, either by using the appliance command line or the IBM MQ Appliance web UI.

Some routine tasks are completed by administrative users on behalf of users, while actions such as changing passwords can be completed by appliance users themselves.

Changing your own password by using the IBM MQ Appliance web UI

You can change your own password by using the IBM MQ Appliance web UI.

Before you begin

You must belong to a group that has the password change permission, or be a privileged user.

About this task

You can change your own password by using the System Control dialog in the IBM MQ Appliance web UI.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Main** > **System Control**. Provided that you have permission to change your own password, the **Change User Password** fields are displayed.
3. Enter your existing password in the **Old Password** field.
4. Enter your new password in the **New Password** field and enter it again in the **Confirm Password** field.
5. Click **Change User Password**.

Changing your own password by using the command line

You can change your own password by using the command line interface.

Before you begin

You must belong to a group that has the password change permission, or be a privileged user.

About this task

You can change your own password by using the password command.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.
2. Log in using your existing password.
3. Type `config` to enter global configuration mode.
4. Type the following command:

```
user-password
```

5. The appliance prompts you to enter the old password, and then to enter and confirm your new password.

Resetting a user's password by using the IBM MQ Appliance web UI

If you are an administrator on the appliance, you can reset the passwords of other users by using the IBM MQ Appliance web UI

Before you begin

You must belong to a group that has the reset password permission, or be a privileged user.

About this task

You use the user account dialog to reset the password for a particular user.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > User Account**.
3. Click on the name of the user whose password you want to reset to open the user account window for that user.
4. From the **Actions** menu, select **Reset password**.
5. Enter the new password in the **New Password** field and then confirm it by typing it again.
6. Click **Reset password**.

Resetting a user's password by using the command line

If you are an administrator on the appliance, you can reset the passwords of other users by using the command line interface.

Before you begin

You must belong to a group that has the reset password permission, or be a privileged user.

About this task

You use the `reset username user new_password` command to reset a user's password.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.
2. Log in as a privileged user.
3. Type `config` to enter global configuration mode.
4. Enter the following command:

```
reset username user new_password
```

Where *user* is the username of the user whose password you are resetting, and *new_password* is the new password.

Forcing a password change by using the IBM MQ Appliance web UI

If you are an administrator on the appliance, you can force a user to reset their password the next time they log in.

Before you begin

You must belong to a group that has the force password change permission, or be a privileged user.

About this task

You use the user account dialog to set the force password change option for a particular user.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > User Account**.
3. Click on the name of the user whose password you want to reset to open the user account window for that user.
4. From the **Actions** menu, select **Force password change**.
5. Confirm that you want to force a password change.

Forcing a password change by using the command line

If you are an administrator on the appliance, you can force a user to reset their password the next time they log in.

Before you begin

You must belong to a group that has the force password change permission, or be a privileged user.

About this task

You use the `user-expire-password` command to force a user to change password.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#).
2. Log in as a privileged user.
3. Type `config` to enter global configuration mode.
4. Enter the following command:

```
user-expire-password user
```

Where *user* is the username of the user for whom you are forcing a password change.

Resetting failed login count by using the IBM MQ Appliance web UI

If you are an administrator on the appliance, you can reset the failed log in count.

Before you begin

You must belong to a group that has the reset failed login counter permission, or be a privileged user.

About this task

You need to reset the failed log in count for a user if they exceed the permitted number of failed login attempts and are locked out of the appliance. You use the user account dialog to reset the count.

Procedure

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Access > User Account**.
3. Click on the name of the user whose failed login count you want to reset.
4. From the **Actions** menu, select **Reset failed login**.
5. Confirm that you want to reset the login count.

Resetting failed login count by using the command line

If you are an administrator on the appliance, you can reset the failed log in count.

Before you begin

You must belong to a group that has the reset failed login counter permission, or be a privileged user.

About this task

You need to reset the failed log in count for a user if they exceed the permitted number of failed login attempts and are locked out of the appliance. You use the reset failed-login command to reset the count.

Procedure

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#).
2. Log in as a privileged user.
3. Type `config` to enter global configuration mode.
4. Enter the following command:

```
reset failed-login user
```

Where *user* is the username of the user whose failed login count you are resetting.

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.

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.

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 `ibmwebspheremqQMgrName`, 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 \(create certificate\)”](#) on page 630.

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”](#) on page 481.

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 488.

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 489.

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 629.

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 `ibmwebspheremqQMgrName`, 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 \(create certificate request\)”](#) on page 632.

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 488.

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”](#) on page 483.

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 489.

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 \(receive certificate\)” on page 641](#).

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 485](#).

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 484](#).
- 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 483](#).

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 491](#).

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 489](#).

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_9.1.0/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 629.

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 \(re-create certificate request\)](#)” on page 642.

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 488.

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 483.

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:

```
detailcert -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 491](#).

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 363](#), and [“Restoring a key repository” on page 364](#).

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 788](#).

Procedure

- To download a file by using the command line interface:
 - a) Connect to the command line of the appliance as described in [“Command line access” on page 207](#).
 - 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 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 on to 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:
 - a) Start the IBM MQ Appliance web UI, and click the **File Management** tab.
 - b) Open the mqpubcert folder.
 - c) 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 `ibmwebspheredmqm1`. The file is downloaded to the home directory on the remote system `192.0.2.2:22`:

```
copy mqpubcert:ibmwebspheredmqm1 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 491](#).

If certificates have been deleted or removed from the `mqpubcert://` folder, you can use the **keybackup** command to copy the full keystore (`.kdb`) file into a tar/archive file in `mqpubcert://` that you can then copy or download, see [“keybackup” on page 637](#). The command also provides you with a password which you subsequently need to use to access the `key.kdb` file. You must copy the `tar.gz` archive file to a machine that has IBM MQ installed, you can then list the `key.kdb` file by using the following command:

```
runmqckm -cert -list -db key.kdb -pw PASSWORD
```

where `PASSWORD` is the password returned by **keybackup**. You can extract from the `key.kdb` file by using the following command:

```
runmqckm -cert -extract -db key.kdb -pw PASSWORD -label LABELNAME
```

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 788](#).

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:
 - Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.
 - Log in as a user in the administrators group.
 - Type the following command to enter configuration mode:

```
config
```

- 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.

- Type **exit** to leave config mode.
- To upload a certificate file from your local system by using the IBM MQ Appliance web UI:
 - Start the IBM MQ Appliance web UI, and click the **File Management** tab.
 - Click **Actions** for the mqpubcert folder.
 - Select **Upload files** from the **Actions** menu.
 - Click **Browse**, and browse for the location of the certificate file on your local system.
 - 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 `ibmwebsphermqm2`. 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:///ibmwebsphermqm2
```

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”](#) on page 491.

Related reference

[“copy”](#) on page 788

This command copies a file to or 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 210).

Procedure

- To upload certificates to your appliance:
 - a) Ensure that you have the following items:
 - A private key to access the appliance certificate.
 - The appliance certificate.
 - Client certificates (optional).
 - b) Connect to the IBM MQ Appliance as described in [“Command line access”](#) on page 207.

- c) Log in as a user in the administrators group.
- d) Type the following command to enter configuration mode:

```
config
```

- e) 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 489.

- To create definition objects for the appliance certificate and key:
 - a) From configuration mode, type `crypto` to enter crypto configuration mode.
 - b) 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
```

- c) Create a crypto certificate definition for the appliance:

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

For example:

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

- d) 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):
 - a) 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
```

- b) 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:
 - a) 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:
 - a) Type `exit` to leave crypto configuration mode.
 - b) Type `write mem` to save your configuration changes.
- To associate the SSL server profile with the web UI:
 - a) From configuration mode, type `web-mgmt` to enter web management configuration mode.
 - b) Enter the following command:

```
ssl-server SSL_Svr_Profile_name
```

For example:

```
ssl-server myappliance1
```

- To save your web management configuration:
 - a) Type `exit` to leave web-mgmt configuration mode.
 - b) Type `write mem` to save your configuration changes.
 - c) Type `exit` again to leave configuration mode.

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

[“idcred” on page 733](#)

This command creates identification credentials.

[“valcred” on page 749](#)

This command enters Validation Credentials mode.

[“key” on page 734](#)

This command creates an alias for a private key.

[“ssl-server” on page 748](#)

This command enters SSL Server Profile mode.

FIPS compliance

Gives a guide to FIPS 140-2 level 1 compliance on the IBM MQ Appliance.

Note: You cannot ensure that all encryption on the appliance is performed using FIPS compliant code paths.

While you can ensure that individual components of the IBM MQ Appliance use FIPS compliant libraries for cryptographic applications, as described in the following sections, there is currently no global way to ensure the system as a whole performs all encryption using only compliant code paths.

Administration interfaces

The appliance has various interfaces that can be used to administer the appliance: SSH, web UI, and REST API. Use the command **crypto-mode-set fips-140-2-11** to tell the appliance administrative process to perform the encryption on these interfaces using a cryptographic software module that is validated to FIPS 140-2 Level 1 (see [“crypto-mode-set” on page 732](#)).

For FIPS compliance and administration interfaces that use MQ Channels (for example, PCF or remote MQSC), see the following section, IBM MQ Channels.

IBM MQ channels

Appliance queue managers can be instructed to use a library that has been tested for FIPS 140-2-l1 compliance for cryptography on all MQ channels. The library is named IBM Crypto for C (ICC). The versions of the library embedded in the IBM MQ Appliance can be displayed using the command `dspmqr -p 64 -v` (see [“dspmqr \(display version information\)” on page 595](#)).

See [Federal Information Processing Standards \(FIPS\) for UNIX, Linux, and Windows](#) in the IBM MQ documentation for more information about IBM MQ channels and FIPS compliance.

IBM MQ clients

For client connections to the appliance, you must ensure that your client is configured for FIPS compliance, see [Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client](#) in the IBM MQ documentation.

Audit trail

The appliance keeps an audit trail in the audit log and the CLI log.

The audit log itself records events such as system configuration changes, successful user log ins, user modifications, and significant system events, such as intrusion detection and firmware information during boot. The CLI log records every CLI command, regardless of whether a configuration change is made. This category provides a record of all commands that are run, including display commands and those within configuration scripts.

V 9.1.4 The information recorded also includes events generated by IBM MQ.

For details about the audit and CLI logs, and how to consume them, see [“Appliance audit logs”](#) on page [512](#).

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.

IBM MQ appliance logs

Log information is generated both by IBM MQ and the appliance itself.

IBM MQ logs

Error and information messages are generated by queue managers as they run. The QMErrorLog service on the appliance writes the log messages to three files in the `mqerr:/qmgrs/queue_manager` directory, where `queue_manager` identifies the queue manager that has generated the logs. The log files are named `AMQERR01.LOG`, `AMQERR02.LOG`, and `AMQERR03.LOG`. The log messages roll over from one file to the next when the file size reaches 32 MB.

mqerr:	Actions...		
qmgrs	Actions...		
QM1	Actions...		
AMQERR01.LOG		15,854	Jul 31, 2018 8:35:19 PM
AMQERR02.LOG		0	Jul 31, 2018 8:35:17 PM
AMQERR03.LOG		0	Jul 31, 2018 8:35:17 PM
QM2	Actions...		
AMQERR01.LOG		15,854	Jul 31, 2018 8:35:25 PM
AMQERR02.LOG		0	Jul 31, 2018 8:35:23 PM
AMQERR03.LOG		0	Jul 31, 2018 8:35:23 PM

Other IBM MQ error and information messages are generated that do not pertain to particular queue managers, but to IBM MQ itself. These messages are written to the file `mqerr:/MQSystem.log`. When the `MQSystem.log` file fills up, its contents are written to a file name `MQSystem.date.time.log`, and a new `MQSystem.log` is started.

mqerr:	Actions...		
mqgrs	Actions...		
<input type="checkbox"/> MQSystem.01.17.18-12.22.38.log		262,207	Jan 17, 2018 5:22:38 PM
<input type="checkbox"/> MQSystem.03.03.18-17.20.31.log		262,207	Mar 3, 2018 10:20:31 PM
<input type="checkbox"/> MQSystem.04.18.18-02.54.58.log		262,255	Apr 18, 2018 2:54:58 AM
<input type="checkbox"/> MQSystem.06.02.18-07.52.51.log		262,207	Jun 2, 2018 7:52:51 AM
<input type="checkbox"/> MQSystem.07.17.18-12.50.44.log		262,207	Jul 17, 2018 12:50:44 PM
<input type="checkbox"/> MQSystem.08.30.18-17.48.31.log		262,188	Aug 30, 2018 5:48:31 PM
<input type="checkbox"/> MQSystem.10.15.18-10.56.08.log		262,171	Oct 15, 2018 10:56:08 AM
<input type="checkbox"/> MQSystem.10.16.17-09.30.55.log		262,151	Oct 16, 2017 2:30:55 PM
<input type="checkbox"/> MQSystem.12.03.17-07.24.38.log		262,224	Dec 3, 2017 12:24:38 PM
<input type="checkbox"/> MQSystem.log		82,279	Oct 29, 2018 11:58:19 AM

First failure data capture (FFDC) files are also generated by IBM MQ on the appliance. FFDC files are written to the `mqerr:/` directory and named `AMQnnnnn.mm.FDC` where `nnnnn` is the ID of the process reporting the error, and `mm` is an identity number that is incremented to ensure that FFDC files always have unique names.

mqerr:	Actions...		
mqgrs	Actions...		
<input type="checkbox"/> AMQ1178672.0.FDC		1,417,199	Sep 17, 2018 9:34:08 PM
<input type="checkbox"/> AMQ1242089.0.FDC		693,681	Sep 17, 2018 9:34:10 PM

You can view the contents of the queue manager logs, system logs, and FFDC files by using the `dspmqr` command, or by viewing them in the IBM MQ Appliance web UI.

Other diagnostic message services are available that you can define to enable you to consume IBM MQ diagnostic information in different ways. You can use these services to generate json logs for both queue managers and IBM MQ system log messages. You can also define a template for logging information to be used when new queue managers are created. When you define these services you can specify filters that control which messages are written to the logs. For more information, see [“IBM MQ diagnostic message services”](#) on page 505.

You can also specify appliance log targets that can collect IBM MQ diagnostic messages. Log targets are the mechanism used to consume general appliance-related error and event messages, and are further described in the following section.

Appliance logs

You can configure log targets on the appliance. Log targets provide an appliance logging facility based on a publish/subscribe model. You can configure multiple log targets, specifying which events are captured and where the log is published to. You can specify that IBM MQ events are captured and published in this way. You can use the IBM MQ Appliance web UI or the command line interface to configure log targets. For more information, see [“Appliance log targets”](#) on page 516.

The appliance also produces error reports. You can generate an error report on demand, or you can configure various automatic reporting options. For example, you can specify that the appliance automatically generates and uploads a first failure data capture file when a failure occurs, or that you are automatically emailed whenever the appliance restarts. For more information, see [Appliance error reports](#).

Audit logs

The appliance generates audit logs so there is always an audit trail to follow. There are two types of audit log:

- The audit log itself records events such as system configuration changes, successful user log ins, user modifications, and significant system events, such as intrusion detection and firmware information during boot. The `audit-log` log file, and its rotations, exist in the `audit:` filesystem. Access to this filesystem is restricted to the CLI and a dedicated section of the web UI (**Status > View Logs > Audit logs**). The audit log cannot be accessed using the standard file management view in the web UI or the REST management interface.
- The CLI log records every CLI command, regardless of whether a configuration change is made. This category provides a record of all commands that are run, including display commands and those within configuration scripts. The web UI and REST management interface are often a front end to the CLI, particularly for configuration changes, so this category also includes log events for commands that are invoked internally to service these interfaces. The `cli-log` log file, and its rotations, exist in the `logtemp:` filesystem, and can be viewed using the standard file management view in the web UI or the REST management interface.

You can configure log targets to subscribe to both the audit log and the CLI log, so that you can filter the logs and write them elsewhere on the appliance, or write them externally to a central syslog server.

V 9.1.4 IBM MQ logs events to both the audit and the CLI logs.

IBM MQ error logs

There are several types of error logs generated by IBM MQ on the 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 [“Configuring queue manager error logs”](#) on page 502.

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 **dspmqrerr** 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 **dspmqrerr** 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:

```
dspmqrerr -s
```

- To list system error log files, enter the following command:

```
dspmqrerr -s -l
```

- To display a specific system error log file, enter the following command:

```
dspmqrerr -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 `MQSystem.18.08.18-9.45.05.log`:

```
dspmqrerr -s MQSystem.18.08.18-9.45.05.log
```

Related reference

[“dspmqrerr \(display mq error log\)” on page 688](#)

Displays the IBM MQ error log files.

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 **dspmqrerr** 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:

```
mqccli
```

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:

```
dspmqrerr -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:

```
dspmqrerr -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:

```
dspmqrerr -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:

```
dspmqrerr -l -m QM1
```

Related reference

[“dspmqrerr \(display mq error log\)” on page 688](#)

Displays the IBM MQ error log files.

Viewing queue manager error log files by using the IBM MQ Appliance web UI

You can view the queue manager error log files by using the IBM MQ Appliance web UI. You can choose to display the most recent error log file, or a specific log file.

About this task

You can view the log file in your browser, and save it to your local computer from your browser, if required.

Procedure

1. Start the IBM MQ Appliance web UI, and click the menu icon  in the title bar.
2. Select **Files** to open the File Management window.
3. Open the `mqerr/qmgrs` folder.
4. Select the log file you want to view. You can view the file, or save it to your local disk, depending on the options offered by your web browser.

Configuring queue manager error logs

You can configure the error logs generated by queue managers so that, for example, certain types of message are filtered out, or messages are only logged once in a specified time period.

About this task

The queue manager error message handling is provided by a service named `QMErrorLog`. By default this service logs messages for each queue manager in the directory `mqerr:/qmgrs/queue_manager` in the files `AMQERR01.LOG`, `AMQERR02.LOG`, and `AMQERR03.LOG`. The log messages roll over from one file to the next when the file size reaches 32 MB.

This base queue manager logging service is controlled by a stanza in the `qm.ini` file of each queue manager. If the `QMErrorLog` service is unmodified, the default settings for the service are used and there is no actual stanza to view in `qm.ini`. You can use commands to create and edit stanzas to customize features such as message filtering and size of log files before they are rotated. You cannot edit the `qm.ini` file directly on the appliance.

Changes to the `QMErrorLog` service stanzas in `qm.ini` only take effect after the queue manager is restarted.

You can, if required, use the diagnostic message services to produce additional, more customized log files, see “[IBM MQ diagnostic message services](#)” on page 505. You can also send queue manager messages to an appliance log target, see “[Appliance log targets](#)” on page 516.

Procedure

- To view the current settings of the `QMErrorLog` service for a particular queue manager, enter the following command:

```
dspmqini -m queue_manager -s QMErrorLog
```

- To configure a stanza for a `QMErrorLog` service, enter the following command:

```
setmqini -m queue_manager -s QMErrorLog -n service_name -k key_name -v key_value
```

The possible options for the `-k` argument are:

- ExcludeMessages
- SuppressMessages
- SuppressInterval

– FileSize

See [“IBM MQ diagnostic message services” on page 505](#) for a fuller description of each key, and possible values.

Viewing the first failure data captures

You can view the first failure data captures (FDCs) by using the **dspmqr** 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:

```
dspmqr -f -l
```

- To display a specific FDC, enter the following command:

```
dspmqr -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:

```
dspmqr -f AMQ12345.FDC
```

Related reference

[“dspmqr \(display mq error log\)” on page 688](#)
Displays the IBM MQ error log files.

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 **dspmqr** command without parameters.

Procedure

- To delete log files by using the command line:
 - a) Enter the IBM MQ administration mode by entering the following command:

```
mqcli
```

b) 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

c) Optionally specify `-y` so that you are not prompted to confirm deletion. For example:

```
dltmqras -a -y
```

- 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**.
- To delete a log file be using the REST API, make the following call:

```
delete https://appliance_address:5554/mgmt/filestore/default/mqerr/filename
```

Related reference

[“dltmqras \(delete mq ras\)” on page 686](#)

Deletes the IBM MQ error log files.

[“dspmqerr \(display mq error log\)” on page 688](#)

Displays the IBM MQ error log files.

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 207](#). 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).

Related reference

[“copy” on page 788](#)

This command copies a file to or from the appliance.

IBM MQ diagnostic message services

The diagnostic message services provides various options for consuming the IBM MQ error logs.

You can define one or more services that configure IBM MQ to output messages to a file. The messages can be filtered by their severity, specific messages can be excluded, and duplicates can be suppressed. Messages can be output in JSON format or as text. You can also specify a service that sends messages to a logging target that you have defined on the appliance (see [Appliance log targets](#)).

You configure the services by using commands that alter stanzas in the `qm.ini` and `mqs.ini` files. (You cannot edit these files directly on the appliance.) The following stanzas are configured in the `qm.ini` and `mqs.ini` files:

DiagnosticSystemMessages

This stanza is defined in `mqs.ini`. It configures a service for IBM MQ diagnostic messages that are not associated with a specific queue manager. These messages equate to those written to the system error log (`MQSystem.log` on the appliance).

DiagnosticMessages

This stanza is defined in `qm.ini`. It configures a service for IBM MQ diagnostic messages that are associated with the queue manager. These messages equate to those written to the queue manager error log (`AMQERR0x.LOG` on the appliance). Changes to this stanza take effect only after the queue manager is restarted.

DiagnosticMessagesTemplate

This stanza is defined in `mqs.ini`. It defines a service for IBM MQ diagnostic messages that is configured for newly created queue managers. When a queue manager is created, this stanza is copied to the new queue manager `qm.ini` file as a `DiagnosticMessages` stanza.

You can specify the following keys (attributes) in the diagnostic message stanzas:

Name

Specifies a unique identifier for the service.

Service

The Service attribute can be set to either `File` for output messages to a file, or `Appliance` for routing messages to log targets. A service of type `Appliance` can be defined once for system messages (`DiagnosticSystemMessages`), once for a template service (`DiagnosticMessagesTemplate`), and once for each queue manager (`DiagnosticMessages`). IBM MQ diagnostic messages are published to all subscribing log targets.

ExcludeMessages

Optionally specify a comma-separated list of message identifiers to exclude from the log.

- 5211 - Maximum property name length exceeded.
- 5973 - Distributed publish/subscribe subscription inhibited
- 5974 - Distributed publish/subscribe publication inhibited
- 6254 - The system could not dynamically load shared library
- 7234 - Number of messages loaded
- 8245 - Entity has insufficient authority to display object
- 9001 - Channel program ended normally
- 9002 - Channel program started
- 9202 - Remote host not available
- 9208 - Error on receive from host
- 9209 - Connection closed
- 9228 - Cannot start channel responder
- 9489 - SVRCONN max instances limit exceeded
- 9490 - SVRCONN max instances per client limit exceeded
- 9508 - Cannot connect to queue manager
- 9524 - Remote queue manager unavailable
- 9528 - User requested closure of channel
- 9545 - Disconnect interval expired
- 9558 - Remote Channel is not available
- 9637 - Channel is lacking a certificate
- 9776 - Channel was blocked by user ID
- 9777 - Channel was blocked by NOACCESS map
- 9782 - Connection was blocked by address

- 9999 - Channel program ended abnormally

SuppressMessages

Optionally specify a comma-separated list of message identifiers that are only output once per time interval. This attribute is not supported in the `DiagnosticSystemMessages` stanza. The following values are valid for this attribute:

- 5211 - Maximum property name length exceeded.
- 5973 - Distributed publish/subscribe subscription inhibited
- 5974 - Distributed publish/subscribe publication inhibited
- 6254 - The system could not dynamically load shared library
- 7234 - Number of messages loaded
- 8245 - Entity has insufficient authority to display object
- 9001 - Channel program ended normally
- 9002 - Channel program started
- 9202 - Remote host not available
- 9208 - Error on receive from host
- 9209 - Connection closed
- 9228 - Cannot start channel responder
- 9489 - SVRCONN max instances limit exceeded
- 9490 - SVRCONN max instances per client limit exceeded
- 9508 - Cannot connect to queue manager
- 9524 - Remote queue manager unavailable
- 9528 - User requested closure of channel
- 9545 - Disconnect interval expired
- 9558 - Remote Channel is not available
- 9637 - Channel is lacking a certificate
- 9776 - Channel was blocked by user ID
- 9777 - Channel was blocked by NOACCESS map
- 9782 - Connection was blocked by address
- 9999 - Channel program ended abnormally

SuppressInterval

Optionally specify the time interval in seconds for the `SuppressMessages` list. The value must be in the range 1 to 86400. The default value is 30. This attribute is not supported in the `DiagnosticSystemMessages` stanza.

Severities

Optionally specify a comma-separated list of severity levels for messages to be output by the message service. The following values are permitted:

- I, Information, or 0
- W, Warning, or 10
- E, Error, or 20 or 30
- S, Stop, or 40
- T, System, or 50
- all

Append a plus character (+) to any of the arguments to specify that severity or higher. For example, "E+" would include messages of severity Error, Stop, and System.

FilePrefix

Optionally specify the prefix for the log file name. The default value is AMQERR.

FileSize

Optionally specify the size at which the log file rolls over (is rotated). The default value is 32 MB.

Format

Optionally specify the format of the log file, which can be `text` or `json`. The file suffix is determined by the type, and is either `.LOG` or `.json`.

Configuring diagnostic message services

You configure the IBM MQ diagnostic message services by using commands to edit the `qm.ini` and `mqs.ini` files.

About this task

Each of the available message services is configured by editing stanzas in the `qm.ini` and `mqs.ini` files. You cannot edit these files directly, so must use the **dspmqini** command to view stanzas and the **setmqini** command to update them. These commands are run at the `mqcli` prompt.

You update the stanzas to configure the `DiagnosticSystemMessages`, `DiagnosticMessages`, and `DiagnosticMessagesTemplate` services. You can also use these methods to update the stanzas for the service that handles the innate IBM MQ error and information messages, **QMErrorLog**, and for the service that makes messages available for log targets, **Appliance**. For the **Appliance** service you can only specify the filtering attributes.

Changes to the `DiagnosticMessages` service stanzas in `qm.ini` only take effect after the queue manager is restarted.

Procedure

- To view the stanza for all the `DiagnosticSystemMessages` services defined in the `mqs.ini` file, enter the following command:

```
dspmqini -g -s DiagnosticSystemMessages
```

- To view the stanza for a specific `DiagnosticSystemMessages` service, enter the following command:

```
dspmqini -g -s DiagnosticSystemMessages -n service_name
```

- To configure a stanza for a `DiagnosticSystemMessages` service, enter the following command:

```
setmqini -g -s DiagnosticSystemMessages -n service_name -k key_name -v key_value
```

The possible options for the `-k` argument are:

- `Service`
- `ExcludeMessages`
- `Severities`
- `FilePrefix`
- `FileSize`
- `Format`

See [“IBM MQ diagnostic message services” on page 505](#) for a fuller description of each key, and possible values.

- To view the stanza for all the `DiagnosticMessages` services defined in the `qm.ini` file for a particular queue manager, enter the following command:

```
dspmqini -m queue_manager -s DiagnosticMessages
```

- To view the stanza for a specific `DiagnosticMessages` service, enter the following command:

```
dspmqini -m queue_manager -s DiagnosticMessages -n service_name
```

- To configure a stanza for a DiagnosticMessages service for a queue manager, enter the following command:

```
setmqini -m queue_manager -s DiagnosticMessages -n service_name -k key_name -v key_value
```

The possible options for the -k argument are:

- ExcludeMessages
- SuppressMessages
- SuppressInterval
- Severities
- FilePrefix
- FileSize
- Format

See [“IBM MQ diagnostic message services” on page 505](#) for a fuller description of each key, and possible values.

- To view the stanza for all the DiagnosticMessagesTemplate services defined in the mqs.ini file, enter the following command:

```
dspmqini -g -s DiagnosticMessagesTemplate
```

- To view the stanza for a specific DiagnosticMessagesTemplate service, enter the following command:

```
dspmqini -g -s DiagnosticMessagesTemplate -n service_name
```

- To configure a stanza for a DiagnosticMessagesTemplate service, enter the following command:

```
setmqini -g -s DiagnosticMessagesTemplate -n service_name -k key_name -v key_value
```

The possible options for the -k argument are:

- ExcludeMessages
- SuppressMessages
- SuppressInterval
- Severities
- FilePrefix
- FileSize
- Format

See [“IBM MQ diagnostic message services” on page 505](#) for a fuller description of each key, and possible values.

Example

The following commands configure a DiagnosticSystemMessages service named SystemErrors that records messages with a severity of error, stop, or system, and writes them to files with the prefix 'SystemErrorsFile'.

```
setmqini -g -s DiagnosticSystemMessages -n SystemErrors -k Service -v File  
setmqini -g -s DiagnosticSystemMessages -n SystemErrors -k Severities -v E+  
setmqini -g -s DiagnosticSystemMessages -n SystemErrors -k FilePrefix -v SystemErrorsFile
```

To view the stanza that you have just configured, enter the following command:

```
dspmqini -g -s DiagnosticSystemMessages -n SystemErrors
```

The following stanza is displayed:

```
DiagnosticSystemMessages:  
Service = File  
FilePrefix = SystemErrorsFile  
Name = SystemErrors  
Severities = E+
```

The following commands configure a DiagnosticMessages service named Filtered that suppresses multiple messages that report the start and end of channel programs within a one minute interval. Messages are written to files with the prefix 'TasksSuppressed'.

```
setmqini -m QM1 -s DiagnosticMessages -n Filtered -k Service -v File  
setmqini -m QM1 -s DiagnosticMessages -n Filtered -k SuppressMessage -v 9001,9002  
setmqini -m QM1 -s DiagnosticMessages -n Filtered -k SuppressInterval -v 60  
setmqini -m QM1 -s DiagnosticMessages -n Filtered -k FilePrefix -v TasksSuppressed
```

To view the stanza that you have just configured, enter the following command:

```
dspmqini -m QM1 -s DiagnosticMessages -n Filtered
```

The following stanza is displayed:

```
DiagnosticMessages:  
Service = File  
SuppressMessage = 9001,9002  
FilePrefix = TasksSuppressed  
SuppressInterval = 60  
Name = Filtered
```

The following example updates the DiagnosticMessages service that makes queue manager log information available to be sent to an appliance log target. The update applies the same filter as applied in the previous example:

```
setmqini -m QM1 -s DiagnosticMessages -n LogTarget -k SuppressMessage -v 9001,9002  
setmqini -m QM1 -s DiagnosticMessages -n LogTarget -k SuppressInterval -v 60  
setmqini -m QM1 -s DiagnosticMessages -n LogTarget -k FilePrefix -v TasksSuppressed
```

To view the stanza that you have just configured, enter the following command:

```
dspmqini -m QM1 -s DiagnosticMessages -n LogTarget
```

The following stanza is displayed:

```
DiagnosticMessages:  
Name = LogTarget  
Service = Appliance  
SuppressMessage = 9001,9002  
FilePrefix = TasksSuppressed  
SuppressInterval = 60
```

For information on how to go on to configure the log target, see [“Appliance log targets” on page 516](#).

Related concepts

[“Default diagnostic message services” on page 511](#)

The appliance has default services configured in the `mqs.ini` file, and in the `qm.ini` file that is created when you create a queue manager.

Default diagnostic message services

The appliance has default services configured in the `mqs.ini` file, and in the `qm.ini` file that is created when you create a queue manager.

Services in the `mqs.ini` file

By default, the `mqs.ini` file contains the following service definitions:

```
DiagnosticSystemMessages:
  Name      = LogTargets
  Service   = Appliance
  Severities = all

DiagnosticMessagesTemplate:
  Name      = LogTargets
  Service   = Appliance
  Severities = all
```

The definition named `LogTargets` makes IBM MQ system messages available to be published to appliance log targets. There can only be one `DiagnosticSystemMessages` stanza with a service of type `Appliance` in the `mqs.ini` file. You can add filtering attributes to this service definition if required (see [“Configuring diagnostic message services”](#) on page 508).

The definition named `DiagnosticMessagesTemplate` ensures that any new queue managers that are created have a `DiagnosticMessages` stanza name `LogTargets` added to their `qm.ini` file that makes diagnostic messages available to be published to log targets.

You can delete the default services from the `mqs.ini` file by using the following commands.

To delete the `LogTargets` service:

```
setmqini -g -s DiagnosticSystemMessages -n LogTargets -d
```

To delete the `DiagnosticMessagesTemplate` service:

```
setmqini -g -s DiagnosticMessagesTemplate -n LogTargets -d
```

You can use the `setmqini` to recreate the two default services if needed.

Services in `qm.ini` files

By default, a `qm.ini` file contains the following service definition:

```
DiagnosticMessages:
  Name      = LogTargets
  Service   = Appliance
  Severities = all
```

The definition named `LogTargets` makes the queue manager messages available to be published to appliance log targets. There can only be one service of type `Appliance` in each `qm.ini` file. You can add filtering attributes to this service definition if required (see [“Configuring diagnostic message services”](#) on page 508).

You can delete the default service from the `qm.ini` file by using the following command:

```
setmqini -m queue_manager -s DiagnosticMessages -n LogTargets -d
```

If you have queue managers that were created in an earlier version of appliance firmware, you might have to manually add the appliance diagnostic message service to your `qm.ini` files. The following command creates the service in a `qm.ini` file:

```
setmqini -m queue_manager -s DiagnosticMessages -n LogTargets -k Service -v Appliance
setmqini -m queue_manager -s DiagnosticMessages -n LogTargets -k Severities -v All
```

Related tasks

[“Configuring diagnostic message services” on page 508](#)

You configure the IBM MQ diagnostic message services by using commands to edit the `qm.ini` and `mqs.ini` files.

Appliance audit logs

The appliance generates audit logs so there is always an audit trail to follow.

There are two logs, the audit log and the CLI log.

Audit log

The audit log records significant appliance events and IBM MQ events.

The appliance logs significant events to the audit log, which is always available to view by using CLI commands (see [“show audit-log” on page 886](#)), or the special category on the web UI (**Status > View Logs > Audit logs**). You can also set up log targets using the audit log category to send this log to another location on the appliance, or to an external central syslog server.

Appliance events

The audit log category keeps a record of significant events on the appliance, including the time of changes made and the identity of the user who made the changes. The log records system configuration changes, successful user log ins, user modifications, and significant system events, such as intrusion detection and firmware information during boot. The log also records a subset of commands that are considered of sufficient significance to be audited. The subset includes the following commands:

- Shutting down or rebooting the appliance (the **shutdown** command).
- The **boot** and **reinitialize** commands, such as use of **boot image** to upgrade the firmware.
- File management commands, including **copy**, **delete**, **move**, **mkdir**, and **rmdir**.
- User management commands, including **create**, **delete**, **modify**, and **password reset**.
- Security access profile configuration.
- Management interface configuration (SSH, REST, web UI).
- SSL and cryptographic configuration, including certificate management.
- Password alias and password map configuration.
- Audit trace configuration.
- NTP and other time settings.
- Error report generation.

These commands are also recorded in the CLI log.

Failed log ins and similar security issues are logged in the auth log category, rather than the audit log category.

Log format for appliance events

Entries in the audit log for appliance events have the following format:

```
Timestamp [Type][Result][Event Code] (User:Domain:Connection Type:IP Address): Event String
```

Timestamp

The time and date that the event occurred.

Type

Specifies one of the following types:

- **conf** – configuration management events
- **dprot** – data protection events
- **file** – filesystem operation events

- **sys** – system operation events
- **user** – user account management events

Result

Reports success or failure.

Event Code

An event code in the format 0x8240nnnn.

User

The user who instigated the event. This field is set to SYSTEM for internal events.

Domain

This field is always set to default on the IBM MQ Appliance.

Connection Type

Specifies one of the following connection types:

- **serial-port** – serial connections
- **secure-shell** – SSH connections
- **web-gui** – web UI connections
- **saml-artifact** – XML management connections (internal use on the MQ Appliance)
- **system** – internal system processes
- **rest-mgmt** – REST management connections

IP Address

The user's IP address, or "*" if the IP address is not known or not applicable.

Event string

The event string is used to provide specific information about the event, such as the command performed, the configuration object that is created or modified or deleted, the user who logged in, and/or the failure reason.

The following example shows the events logged when creating a new user using the web UI. Sensitive values, like passwords, are redacted in the log by substituting the string "(*****)":

```
20190828T085452.114Z [user][success][0x8240001f] (admin:default:web-gui:9.174.20.193): (config
user joanna)# user joanna
20190828T085452.114Z [conf][success][0x8240001f] (admin:default:web-gui:9.174.20.193): (config
user joanna)# reset
20190828T085452.115Z [user][success][0x8240001f] (admin:default:web-gui:9.174.20.193): (config
user joanna)# password "(*****)"
20190828T085452.115Z [user][success][0x8240001f] (admin:default:web-gui:9.174.20.193): (config
user joanna)# suppress-password-change
20190828T085452.115Z [user][success][0x8240001f] (admin:default:web-gui:9.174.20.193): (config
user joanna)# access-level group-defined
20190828T085452.115Z [user][success][0x8240001f] (admin:default:web-gui:9.174.20.193): (config
user joanna)# group MQAdmin
20190828T085452.159Z [conf][success][0x8240001c] (admin:default:web-gui:9.174.20.193): user
'joanna' - Configuration added
20190828T085452.164Z [conf][success][0x8240001c] (admin:default:web-gui:9.174.20.193): user
'joanna' - Configuration settings applied
```

IBM MQ events

▶ V 9.1.4

Significant IBM MQ events are also added to the audit log. Audit events are generated for all IBM MQ control command operations that perform an action or modify the configuration (that is, all operations besides those that simply display data).

Log format for IBM MQ events

▶ V 9.1.4

The audit log entries have the same basic format as the appliance audit log entries, but some of the fields are used differently:

Type

Gives one of the following types:

- **conf** – most events are configuration management events
- **user** – events related to messaging users

Event Code

Always set to 0x8240001f for IBM MQ events.

Event string

The event string is the IBM MQ command performed and its parameters.

The following example shows the events logged when creating a high availability queue manager and a new IBM MQ user . Sensitive values, like passwords, are redacted in the log by substituting the string "(*****)":

```
20190828T102741.361Z [conf][success][0x8240001f] (admin:default:secure-shell:9.20.33.218):
(mqcli)# crtmqm -sx QM1
20190828T102741.361Z [user][success][0x8240001f] (admin:default:secure-shell:9.20.33.218):
(mqcli)# usercreate -u jamie -g mqm -p "(*****)"
```

CLI log

The CLI log creates a record of all CLI commands run on the appliance, including `mqcli` commands.

CLI events record the input of commands, not the result or outcome. CLI trace records are generated when a command is issued, irrespective of whether the command is included in the audit trace. The web UI and REST management interface are often a front end to the CLI, particularly for configuration changes, so this category also includes log events for commands that are invoked internally to service these interfaces.

Appliance CLI

CLI log events have the following format:

```
1,Timestamp,Domain,Category,Priority,,,0,,,Event Code,,, "(User:Domain:Connection Type:IP Address): Command"
```

If a command is executed while processing a configuration script using **exec** (for example, when replaying the saved configuration during boot) then the log events have the following format:

```
1,Timestamp,Domain,Category,Priority,,,0,,,Event Code,,, "(User:Domain:Connection Type:IP Address:Script:Line): Command"
```

Timestamp

The time and date that the event occurred.

Domain

This field is always set to default on the IBM MQ Appliance.

Category

Always set to CLI.

Priority

Always set to debug.

Event code

An event code in the format 0x8100nnnn.

User

The user running the command. Is set to admin for internal events.

Domain

This field is always set to default on the IBM MQ Appliance.

Connection Type

Specifies one of the following connection types:

- **serial-port** – serial connections
- **secure-shell** – SSH connections
- **web-gui** – web UI connections
- **saml-artifact** – XML management connections (internal use on the MQ Appliance)
- **system** – internal system processes
- **rest-mgmt** – REST management connections

IP Address

The user's IP address, or "*" if the IP address is not known or not applicable.

Script

If applicable, the script that the command was called from.

Line

If applicable, the line of the script that the command was called from.

Command

The command.

The following example shows the events logged when creating a new user using the web UI. Sensitive values, like passwords, are redacted in the log by substituting the string "(*****)":

```
1,20190828T085452.113Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): top"
1,20190828T085452.114Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): configure terminal"
1,20190828T085452.114Z,default,cli,debug,,0,,0x810001f1,,,"(admin:default:web-gui:9.174.20.193): Test: %if% available ""user"" true; continuing"
1,20190828T085452.114Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): user joanna"
1,20190828T085452.114Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): reset"
1,20190828T085452.115Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): password ""(*****)""""
1,20190828T085452.115Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): suppress-password-change"
1,20190828T085452.115Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): access-level group-defined"
1,20190828T085452.115Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): group MQAdmin"
1,20190828T085452.164Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:web-gui:9.174.20.193): exit"
```

IBM MQ CLI

V 9.14 CLI log events for `mqcli` commands have the same format as the appliance CLI log events. Entry and exit entries are recorded to show entry to and exit from the `mqcli` subshell:

```
1,20190828T101504.349Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:secure-shell:9.20.33.218): mqcli"
1,20190828T101508.763Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:secure-shell:9.20.33.218): exit"
```

The `mqcli` command processor generates a CLI trace event for every recognized command. No trace event is generated for unrecognized commands. Sensitive values in command parameters, such as passwords and secrets, are redacted using the existing convention, which replaces each sensitive value with "(*****)".

The following example shows CLI log entries for commands that create a new high availability queue manager and a new IBM MQ user.

```
1,20190828T102741.361Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:secure-shell:9.20.33.218): crtmqm -sx QM1"
1,20190828T102741.361Z,default,cli,debug,,0,,0x810001f0,,,"(admin:default:secure-shell:9.20.33.218): usercreate -u jamie -g mqm -p ""(*****)""""
```

Writing audit log data to a syslog target

You can write audit log data to a syslog target, as shown in the following example.

To configure a syslog log target:

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Miscellaneous > Manage log targets**.
3. Click **New**.
4. Specify a name for the log target, for example, "Logstash_target".
5. In the **General Configuration** section, complete the following steps:
 - a. Specify a **Target Type** of **Syslog** to send log events over UDP, or select **Syslog-tcp** to send log events over TCP.
 - b. Select the **user** syslog facility.
6. In the **Source Configuration** section, complete the following steps:
 - a. Specify the IP address of the interface over which the syslog events are transmitted to the remote recipient.
 - b. Specify a **Local Identifier** to help you identify log entries sent to syslog. For example, you could specify the system name of the appliance.
7. In the **Destination Configuration** section, specify the IP address of the target syslog server, and optionally specify the port number (if you do not specify a port number, the default UDP port 514 or the default TCP port 25 is used).
8. Scroll down and open the **Event Subscriptions** section and complete the following steps:
 - a. Click **Add**.
 - b. Select an **Event Category** of **audit**.

If you want to write the CLI log to the syslog target too, add another event subscription, and choose an **Event Category** of **cli**.
9. Click **Apply**, and when the window closes, click **Save changes**.

Related concepts

[“Appliance log targets” on page 516](#)

Log targets provide an appliance logging facility based on a publish/subscribe model.

[“Log target examples” on page 520](#)

These examples show how to define log targets to capture IBM MQ event information.

Appliance log targets

Log targets provide an appliance logging facility based on a publish/subscribe model.

You can define one or more log targets. Each target subscribes to a set of system events based on their origin and severity. Log targets capture events that occur because of some internal process or hardware status change. For example, you could configure a log target that captured network events with a severity of 'error'. An appliance supports a maximum of 500 log targets.

An IBM MQ diagnostic message service named `Appliance` automatically makes IBM MQ errors and messages available to be added to log targets. (You can define filtering attributes for the `Appliance` service, as described for other diagnostic message services in [“IBM MQ diagnostic message services” on page 505](#), but it is easier to define filters by using the log target tools in the IBM MQ Appliance web UI.)

You can define where the log information is written for each target and the format that it is written in. You can specify the following destination for a log:

Console

Writes log entries to the screen with Telnet, SSH, or command-line access through the serial port.

Cache

Writes log entries to system memory.

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.

SMTP

Forwards log entries as email to the configured remote SMTP servers and email addresses. Before sending, the contents of the log can be encrypted or signed. The processing rate can be limited.

File

Writes log entries to a file on the appliance. The file can be archived using the rotate or upload method. The upload method uses FTP, SFTP, or SMTP. The file can be sent as email. The entire file can be encrypted and signed.

SOAP

Forwards log entries as SOAP messages. The URL can be set. The processing rate can be limited. Select a remote SOAP target.

SNMP

Forwards log entries as SNMP traps to configured recipients. The processing rate can be limited.

NFS

Writes log entries to a file on a remote NFS server. The file can be archived using the rotate method. The processing rate can be limited.

Each log target can be configured to output events in a specific format, which can be configured dependent upon the log target type. The supported formats are:

Text

Events are written as formatted text.

Raw

Events are written as unformatted text.

XML

Events are written as XML.

CBE

Events are written in IBM common base event (ICBE) format.

CSV

Events are written in comma-separated value (CSV) format.

Each log target can be configured to use the following timestamp formats:

syslog

Uses the syslog-style timestamp.

numeric

Uses a numeric timestamp.

zulu

Milliseconds used as the timestamp format in Coordinated Universal Time (UTC) in log messages `yyyymmddThhmmss+00:00`.

Configuring log targets using the IBM MQ Appliance web UI

You can configure a logging target by using the IBM MQ Appliance web UI.

About this task

First you configure the log target to specify where and how it is written. You then specify a subscription to one or more events for your log target. You can restrict the messages in log targets by object filters, event

category, and event priority. You could, for example, subscribe to **all** events and then specify object, event, or IP address filters to extract the exact information that you are interested in to your log target.

If you want to set up a log target to capture IBM MQ diagnostic messages, you set an event category of `qmgr`. You can filter the log to include or exclude messages with event codes that correspond to IBM MQ message codes. You can specify an object type and an object name to define a log target for messages from a particular queue manager.

Procedure

To configure a log target:

1. Start the IBM MQ Appliance, and click the Administration icon.
2. Select **Miscellaneous > Manage Log Targets**.
3. Click **New** to create a new log target and enter a name for it. Optionally use the **Comments** field to describe the log target that you are defining.
4. Choose a **Target Type**. This is where the log is written to, available types are described in [“Appliance log targets”](#) on page 516.
5. Choose a **Log Format**. This is the format that the log is written in, available formats are described in [“Appliance log targets”](#) on page 516.
6. Choose a **Timestamp Format** to specify what format of timestamp are added to log entries. Available formats are described in [“Appliance log targets”](#) on page 516.
7. Specify a subscription for the log target:
 - a) Scroll down to **Event Subscriptions**, open it and click **Add**.
 - b) Select an **Event Category** from the list. You can select **all** if required, and then narrow down the messages actually written to the log target by specifying filters. If you want to define a log target for IBM MQ diagnostic messages, select the category `qmgr`.
 - c) Specify the minimum priority of an event for that event to be reported. For example, if you choose a **Minimum Event Priority** of **error**, then all events with priority **error**, **critical**, **alert**, and **emergency** are logged.
 - d) If required, click **Add** again and add further combinations of event category and minimum priority.
8. Optionally set up a filter that includes, or excludes, events that have a particular event code.
 - a) Open the **Event Filters** section.
 - b) To set up a filter that includes events, under **Event Subscription Filter**, click **Add** if you know the code of the event that you want to add and type it in, or click **Select Code** to choose a code from a list. Specify or select more codes to add more events to your filter. Event codes for IBM MQ diagnostic messages have the form `0x8d00amq_errorcode`. For example, the error AMQ9002I is identified by the event code `0x8d009002`.
 - c) To set up a filter that excludes events, under **Event Suppression Filter**, click **Add** if you know the code of the event that you want to add and type it in, or click **Select Code** to choose a code from a list. Specify or select more codes to add more events to your filter.
9. Optionally set up a filter that includes messages generated by classes of object, or by a particular instance of an object. For example, you could specify that messages generated about any of the appliance's Ethernet interfaces are written to the log, or that only message generated about interface `eth20` are written. If you are defining a logging target for IBM MQ diagnostic messages, you can set **Object Type** to `qmgr` and **Object Name** to the name of a queue manager to receive messages from the specified queue manager.
 - a) Open the **Object Filters** section.
 - b) Select the type of object that you want to receive messages from the **Object Type** list.
 - c) Optionally, specify an instance of that object type in the **Object Name** field to narrow down the source of messages.
 - d) Optionally, select the **Add Referenced Objects** option to include messages from objects referenced by the object that you have selected.

- e) Click **Add** to add another object filter.
10. Optionally, set up a filter that includes messages associated with a specified IP address. If you are defining a log target for IBM MQ diagnostic messages, you can specify an IP address to filter IBM MQ messages by remote end-point (for example, identifying a client).
 - a) Open the **IP Address Filters** section.
 - b) Click **Add** and specify the IP address in the **IP Address** field.
 - c) Click **Add** to specify a further IP address filter.
 11. Optionally, specify one or more event triggers. Event triggers allow you to trigger an action on the appliance when a particular event occurs. A typical use might be to generate an error report when a rarely observed but recurring message is logged. When an event is triggered, the triggering message is logged and the next entry in the log is the triggered command. This second log entry is prefaced with the words "Event Trigger command". To track event triggers, configure a log target that subscribes only to the message ID of the triggered commands. The triggered commands are logged at the "notice" level, so ensure that you set the log level to "notice".
 - a) Open the **Event Triggers** section and click **Add**.
 - b) Specify the msgid of the message that you want to trigger your action in the **Message ID** field (you can retrieve msgids by viewing existing logs).
 - c) Optionally, specify a regular expression. This is used to further filter the messages that trigger the event. You enter a regular expression to identify text that the message should contain before the event is triggered.
 - d) Select **Only Once** to specify that the action is triggered only the first time that a matching message is logged. Clear **Only Once** to perform the action every time that a matching message is logged.
 - e) Select **Only this Trigger** to specify that only the commands specified by this trigger event are run when a matching message is logged. Any subsequent commands that might be triggered by a message with that id are ignored.
 - f) In the **CLI Command** field, enter the command that is to be run when the triggering message is logged.
 - g) Click **Add** to specify another event trigger.

Configuring log targets using the command line

You can configure a logging target by using the IBM MQ Appliance command line interface.

About this task

First you configure the log target to specify where and how it is written. You then specify a subscription to one or more events for your log target. You can restrict the messages in log targets by object filters, event category, and event priority. You could, for example, subscribe to all events and then specify object, event, or IP address filters to extract the exact information that you are interested in to your log target.

Procedure

To configure a log target:

1. Connect to the IBM MQ Appliance as described in [“Command line access” on page 207](#). 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`, capturing diagnostic messages from a queue manager named `QM1`:

```

mqa(config)# logging target QM1_messages
New Log Target configuration

mqa(config logging target QM1_messages)# summary "Remote logging to rmach.hursley.ibm.com
for queue manager QM1"
mqa(config logging target QM1_messages)# type syslog-tcp
mqa(config logging target QM1_messages)# timestamp syslog
mqa(config logging target QM1_messages)# local-ident "warrior12"
mqa(config logging target QM1_messages)# remote-address "198.51.100.0" "1514"
mqa(config logging target QM1_messages)# local-address 198.51.100.12
mqa(config logging target QM1_messages)# event "qmgr"
mqa(config logging target QM1_messages)# object qmgr QM1
mqa(config logging target QM1_messages)# exit
mqa(config)# write mem

```

Related concepts

[“Log target commands” on page 855](#)

Log target mode provides the commands to create or modify a log target.

Log target examples

These examples show how to define log targets to capture IBM MQ event information.

Writing IBM MQ log messages to a file log target

In this example both information generated by queue managers plus IBM MQ system information is written to a file on the appliance. The default appliance services in `mq.s.ini` and each `qm.ini` file have not been altered. The log target is configured using the IBM MQ Appliance web UI. (If you have queue managers that were created in an earlier version of appliance firmware, you might have to manually add the appliance diagnostic message service to your `qm.ini` files, see [“Default diagnostic message services” on page 511.](#))

To configure the log target:

1. Start the IBM MQ Appliance web UI and click the Administration icon .
2. Select **Miscellaneous > Manage log targets**.
3. Click **New**.
4. Specify a name for the log target, for example, "MQ_example_target".
5. In the **General Configuration** section, complete the following steps:
 - a. Specify a **Target Type** of **File**.
 - b. Ensure that the selected **Log Format** is **XML**.
6. In the **Destination Configuration** section, specify a **File Name** in the form `logstore:///filename`. All log target files are stored in the Logstore directory, and the field will only accept an entry in that form.
7. Scroll down and open the **Event Subscriptions** section and complete the following steps:
 - a. Click **Add**.
 - b. Select an **Event Category** of **qmgr**.
8. Click **Apply**, and when the window closes, click **Save changes**.

Create some activity on your queue managers to create some log entries.

To view the log target:

- In the IBM MQ Appliance web UI, click the Status icon .
- Select **View Logs > Status Logs**.
- Select your log target from the **Target** list.

Your log entries are displayed.

Writing log messages for a specific queue manager

In this example, log messages generated by a specific queue manager are written to a file on the appliance.

To configure the log target:

1. Start the IBM MQ Appliance web UI and click the Administration icon.
2. Select **Miscellaneous > Manage log targets**
3. Click **New**.
4. Specify a name for the log target, for example, "QM1_target".
5. In the **General Configuration** section, complete the following steps:
 - a. Specify a **Target Type** of **File**.
 - b. Ensure that the selected **Log Format** is **XML**.
6. In the **Destination Configuration** section, specify a **File Name** in the form `logstore:///filename`. All log target files are stored in the `Logstore` directory, and the field will only accept an entry in that form.
7. Scroll down and open the **Event Subscriptions** section and complete the following steps:
 - a. Click **Add**.
 - b. Select an **Event Category** of **qmgr**.
8. Scroll up and open the Object Filters section and complete the following steps:
 - a. Click **Add**.
 - b. Select an **Object Type** of **MQ Queue Manager**.
 - c. Specify the name of the queue manager in the **Object Name** field.
9. Click **Apply**, and when the window close, click **Save changes**.

You can view the log by selecting **Status > View Logs > Status Logs** as described for the first example.

Writing log messages to a syslog target

The previous examples showed how to write IBM MQ event data to files in XML-format. You can also write event data to a syslog target, as shown in the following example. In this example both information generated by queue managers plus IBM MQ system information is written to the syslog target.

To configure a syslog log target:

1. Start the IBM MQ Appliance web UI and click the Administration icon.
2. Select **Miscellaneous > Manage log targets**
3. Click **New**.
4. Specify a name for the log target, for example, "Logstash_target".
5. In the **General Configuration** section, complete the following steps:
 - a. Specify a **Target Type** of **Syslog** to send log events over UDP, or select **Syslog-tcp** to send log events over TCP.
 - b. Select the **user** syslog facility.
6. In the **Source Configuration** section, complete the following steps:
 - a. Specify the IP address of the interface over which the syslog events are transmitted to the remote recipient.
 - b. Specify a **Local Identifier** to help you identify log entries sent to syslog. For example, you could specify the system name of the appliance.
7. In the **Destination Configuration** section, specify the IP address of the target syslog server, and optionally specify the port number (if you do not specify a port number, the default UDP port 514 or the default TCP port 25 is used).

8. Scroll down and open the **Event Subscriptions** section and complete the following steps:
 - a. Click **Add**.
 - b. Select an **Event Category** of **qmgr**.
9. Click **Apply**, and when the window close, click **Save changes**.

When you select a syslog target you can create various configurations to consume and examine your IBM MQ log information. An example of deploying logs to Logstash for integration with Elasticsearch and Kibana is documented at <https://github.com/ibm-messaging/mq-appliance/tree/master/elastic>.

Related concepts

[“Writing audit log data to a syslog target” on page 516](#)

You can write audit log data to a syslog target, as shown in the following example.

Monitoring system resource usage

You can monitor the use of system resources by 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
- CPU average load (in 1, 5, and 15 minute averages)
- 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

- 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%]
```

Related tasks

[“Monitoring system resource usage by using the amqsrua command” on page 523](#)

You can use the **amqsrua** command to query metadata that is related to the system resource usage of a queue manager.

Related reference

[“status” on page 881](#)

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.

[“show” on page 715](#)

This command displays configuration or status information

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)#
```

Related tasks

[“Monitoring system resource usage by using the status command” on page 522](#)

You can monitor the use of system resources on the appliance by using the **status** command on the command line.

Monitoring system resource usage by using the IBM MQ Console

You use the Charts widget in the IBM MQ Console to view monitoring data for queue managers.

About this task

You add a Charts widget to your dashboard and then configure it to monitor a particular aspect of resource usage. You can create many instances of the Charts widget to display different data. The data is displayed in a chart format.

Data is collected at 10-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 must have at least one running queue manager before you can configure a chart widget.

Procedure

1. Add a Charts widget to your dashboard:

- a) Click **Add widget**  **Add widget**.
- b) Select **Charts**.

2. Configure the Charts widget to show data:

- a) Click the configure icon  in the title bar of the Charts widget.
- b) Optional: Enter a **Widget title**. This title is shown in the title bar of the widget.
- c) Select the **Resource class** to monitor:

Platform central processing units

Monitor the usage of the CPUs.

Platform persistent data stores

Monitor the use of disk resource.

API usage statistics

Monitor API calls.

API per-queue usage statistics

Monitor API calls by individual queues. When you choose this class, you specify the queue name to monitor in the **Object** field.

d) Select the **Resource type** to monitor.

The resource types that are available to select depend on the resource class that is selected. The following table shows the 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.

<i>Table 44. Resource types (continued)</i>		
Class	Type	Description
	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.
	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 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 sync points by the queue manager.
	Subscribe	Select this type to view data that is related to MQSUB calls.
	Publish	Select this type to view data about published messages.

<i>Table 44. Resource types (continued)</i>		
Class	Type	Description
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.

e) Select the **Resource element** to monitor:

The resource elements that are available to select depend on the resource class and resource type that are selected. The following tables show the resource elements:

<i>Table 45. Elements for Platform central processing units resources</i>		
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 load average over 1 minute.
	CPU load - five-minute average	Shows the load average over 5 minutes.
	CPU load - fifteen-minute average	Shows the load average over fifteen minutes.
CPU performance - running queue manager	RAM free percentage	Shows the percentage of free RAM memory.
	RAM total bytes	Shows the total bytes of RAM configured.
	RAM total bytes - estimate for queue managers	Estimates the total bytes of RAM in use by the queue managers that are being monitored.
CPU performance - running queue manager	User CPU time - percentage estimate for queue manager	Estimates the percentage of CPU use in user state for processes that are 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 that are 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.

<i>Table 46. Elements for Platform persistent data stores resources</i>		
Type	Element	Description
Disk usage - platform wide	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 free.
	MQ FDC file count	Shows the current number of FDC files.
	Appliance data - bytes in use	Shows the overall disk usage.
	Appliance data - free space	Shows the overall free space.
	System volume - bytes in use	
	System volume - free space	
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 that is 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 is 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.

<i>Table 46. Elements for Platform persistent data stores resources (continued)</i>		
Type	Element	Description
	Log - write latency	Shows a measure of the latency when writing synchronously to the queue manager recovery log.

<i>Table 47. Elements for API usage statistics resources</i>		
Type	Element	Description
MQCONN and MQDISC	MQCONN/MQCONN count	Shows the number of calls to MQCONN and MQCONN.
	Failed MQCONN/MQCONN count	Shows the number of failed calls to MQCONN and MQCONN.
	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.

<i>Table 47. Elements for API usage statistics resources (continued)</i>		
Type	Element	Description
	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 that are removed from queues by MQGET.
	Interval total destructive get - byte count	Bytes of data that is removed from queues by MQGET.
	Non-persistent message destructive get - count	Number of non-persistent messages that are removed from queues by MQGET.
	Persistent message destructive get - count	Number of persistent messages that are 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 that are returned to MQGET.
	Got persistent messages - byte count	Shows a count of bytes of persistent messages that are returned to MQGET.

<i>Table 47. Elements for API usage statistics resources (continued)</i>		
Type	Element	Description
	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.

<i>Table 47. Elements for API usage statistics resources (continued)</i>		
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.

<i>Table 47. Elements for API usage statistics resources (continued)</i>		
Type	Element	Description
Publish	Topic MQPUT/MQPUT1 interval total	The number of messages that are 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 that are published to subscribers.
	Published to subscribers - byte count	Shows the byte count of messages that are published to subscribers.
	Non-persistent - topic MQPUT/MQPUT1 count	Shows the number of non-persistent messages that are put to topics.
	Persistent - topic MQPUT/MQPUT1 count	Shows the number of persistent messages that are put to topics.
	Failed topic MQPUT/MQPUT1 count	Shows the number of failed attempts to put to a topic.

<i>Table 48. Elements for API per-queue usage statistics resources</i>		
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.

Table 48. Elements for API per-queue usage statistics resources (continued)

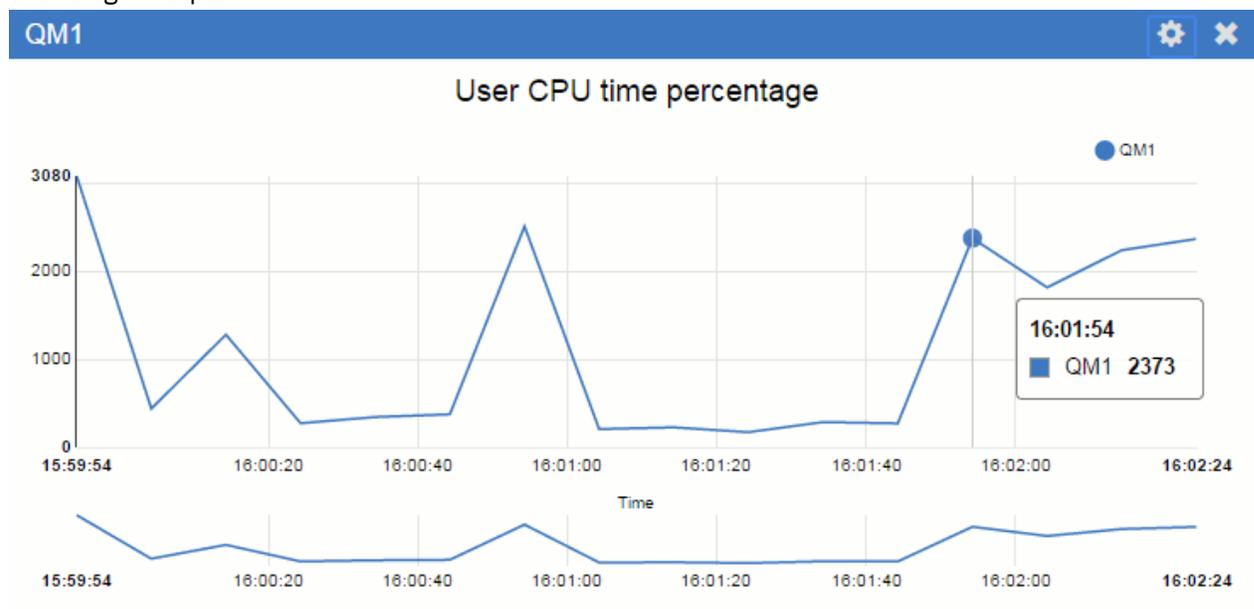
Type	Element	Description
	MQPUT1 non-persistent message count	Shows the number of non-persistent messages that are put by MQPUT1.
	MQPUT1 persistent message count	Shows the number of persistent messages that are put by MQPUT1.
	Non-persistent byte count	Shows the number of bytes put in non-persistent 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 that are removed from the queue by MQGET.
	Destructive MQGET persistent message count	Number of persistent messages that are removed from the queue by MQGET.
	Destructive MQGET non-persistent byte count	Shows a count of bytes of non-persistent messages that are returned to MQGET.
	Destructive MQGET persistent byte count	Shows a count of bytes of persistent messages that are 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.

Table 48. Elements for API per-queue usage statistics resources (continued)		
Type	Element	Description
	Queue purged count	Shows a count of queues that have been purged.
	Average queue time	
	Queue time	

- f) 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.
- g) Click **Save**.

Results

After you configure the widget, there is a short delay before data is displayed in the chart. Data is displayed along a time axis. Each data point represents the end of the 10-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:



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 888	Displays the current time and appliance uptime.
“show file” on page 891	Displays a specified printable file.
“show firmware-version” on page 892	Displays the current firmware version, without image type and installation date.
“show firmware” on page 891	Displays the current firmware version, with image type and installation date.
“show ipaddress” on page 892	Provides IP address information about interfaces.
“show link-aggregation-member-status” on page 893	Lists members in link aggregation interfaces.
“show link-aggregation-status” on page 894	Provides statistics for aggregate interfaces.
“show link” on page 895	Provides status about all interfaces on the appliance.
“show load” on page 895	Displays task level system usage.
“show log” on page 896	Displays the appliance default log.
“show logging” on page 896	Displays a specified appliance log.
“show loglevel” on page 897	Displays the log-level for logging targets.
“show ndcache” on page 898	
“show network-interface” on page 899	Shows generic status of all network interfaces on the appliance.
“show ntp-refresh” on page 900	Lists the refresh status for the current NTP server.
“show raid-array” on page 901	Displays the status of the RAID array.
“show raid-battery-module” on page 901	Displays the information about the battery backup unit of the RAID controller.
“show raid-logical-drive” on page 902	Displays the status of the RAID logical drive.
“show raid-physical-drive” on page 903	Displays the status of the RAID physical drive.
“show raid-ssd” on page 904	Displays the estimated remaining life of the solid state disks
“show route” on page 905	Shows the routing table.
“show sensors-current” on page 905	Displays the values for sensors that read electrical current.
“show sensors-fans” on page 905	Displays the values for sensors that read the speed of the fans.
“show sensors-other” on page 906	Displays the status of sensors that have true or false values.
“show sensors-temperature” on page 907	Displays the values for sensors that read temperatures.
“show sensors-voltage” on page 907	Displays the values for sensors that read voltage.
“show services” on page 908	

Command	Description
“show services-memory” on page 908	
“show system” on page 909	Displays the System Settings configuration.
“show tcp-connections” on page 909	Lists the number of TCP connections in specific states.
“show tcp-table” on page 909	Lists the current TCP connections.
“show tcp” on page 909	Lists the current TCP connections followed by the number of connections in each state.
“show throughput” on page 910	Displays interface-specific traffic statistics.
“show time” on page 910	Displays the current time and appliance uptime.
“show users” on page 911	Lists all users who are currently logged in to the appliance.
“show version” on page 911	Displays the version of the firmware and libraries.

Related concepts

[Arguments for show command](#)

Use the show command to view the status of the appliance or an aspect of the appliance configuration.

Related reference

[Show command](#)

This command displays configuration or status information

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 523](#).

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 349.

Monitoring queue manager health

You have a choice of interfaces you can use to monitor the health of your queue managers.

The IBM MQ Appliance gives access to machine status through a mechanism called status providers. These status providers can be accessed by using the CLI, REST API, and IBM MQ Appliance web UI interfaces. There is a “QueueManagersStatus” status provider, which you can use to view the running status of all your queue managers on the appliance, including high availability and disaster recovery information.

Using the CLI

To monitor the health of queue managers by using the command line interface, use the **show** command:

```
show qm-status
```

The appliance displays the status of the queue managers that are currently running, for example:

```
mqa# show qm-status

Name Status          CPU usage Used memory Used File System Total File System HA Role HA
Status DR Role DR Status
-----
sam1 Ended immediately 0          0          117          3072          NA
NA      Primary Normal
```

Using the REST API

You can retrieve the status of a queue manager by constructing a URI that queries the status provider. See [“Monitoring the appliance by using the REST management interface”](#) on page 553 for more information on constructing the URI.

For example, the following request:

```
GET 'https://example.com:5554/mgmt/status/default/QueueManagersStatus'
```

might result in the following response:

```
{
  "_links": {
    ...
  },
  "QueueManagersStatus": {
    "Name": "sam1",
    "Status": "Ended immediately",
    "CpuUsage": 0,
    "UsedMemory": 0,
    "UsedFs": 117,
    "TotalFs": 3072,
    "HaRole": "NA",
    "HaStatus": "NA",
    "DrRole": "Primary",
    "DrStatus": "Normal"
  }
}
```

IBM MQ Appliance web UI

You can use the IBM MQ Appliance web UI to query the QueueManagersStatus status provider:

1. Start the IBM MQ Appliance web UI, and click the status icon .
2. Select **MQ > Queue Managers Status**.

The status of queue managers running on the appliance is displayed, for example:

Queue Managers Status

Name	Status	CPU usage (%)	Used memory (MB)	Used File System (MB)	Total File System (MB)	HA Role	HA Status	DR Role	DR Status
sam2	Ended immediately	0	0	117	3,072	NA	NA	Primary	Normal

Monitoring the health of high availability and disaster recovery configurations

Status notifications are issued when HA and DR queue managers, and the HA group, changes state. You can use these notifications to track the health of the HA or DR configuration.

These notifications are written as messages to `mqserr:/MQSystem.log` (see [“IBM MQ appliance logs”](#) on page 497), providing an audit of HA and DR events. Both primary and secondary instances of HA and DR queue managers issue notifications.

You can set up appliance log targets and send these HA and DR health notifications to a log target to provide real-time monitoring of HA and DR configurations (see [“Appliance log targets”](#) on page 516).

The following topics give examples of sequences of notifications that can be issued by HA or DR configurations.

HA notification examples

This topic illustrates common sequences of notifications that might be issued by a High Availability (HA) configuration.

Creating an HA queue manager

When you create a new HA queue manager, a series of notifications are issued. First you issue the command to create the queue manager:

```
crtmqm -sx QM_HA
```

As the creation progresses the following sequence of state notifications are seen.

Replication notifications

```
AMQ3592E: HA status for queue manager QM_HA is 'Inconsistent'
```

As the data replication resources are created for the queue manager on the remote appliance, this state indicates it has yet to receive data from the local appliance.

```
AMQ3598W: HA status for queue manager QM_HA is 'Synchronization in progress'
```

After the data replication resources on both appliances connect they begin synchronizing the data from the local appliance to the remote appliance. If replication is interrupted, for example, by interruption of the network, the state returns to Inconsistent.

```
AMQ3599I: HA status for queue manager QM_HA is 'Normal'
```

After the secondary appliance has synchronized all the data from the primary appliance, all disk writes on the primary appliance are synchronously replicated to the secondary appliance. This is the Normal state.

High Availability notifications

Alongside the notifications relating to the health of the data replication are additional notifications relating to the high availability state.

```
AMQ3912I: HA preferred appliance for queue manager QM_HA is 'Appliance1'
```

This message is issued on both appliances to indicate which appliance has been selected as the preferred location for the queue manager.

The appliance on which **crtmqm** is issued is automatically selected as the preferred location, it can be changed by issuing the **sethapreferred** command on the remote appliance.

```
AMQ3902I: HA role for queue manager QM_HA is primary
```

This message is seen on the appliance where **crtmqm** is issued, as it promotes the local instance of the queue manager to the primary role.

```
AMQ3903I: HA role for queue manager QM_HA is secondary
```

The remote appliance reports it is taking on the secondary role for the queue manager being created.

```
AMQ3913I: HA control for queue manager QM_HA is enabled
```

The final notification triggered by the **crtmqm** command indicates the queue manager has been put under control of the HA subsystem, which then monitors the two appliances to determine where best to run the queue manager.

Shortly after this message is issued the queue manager automatically starts on the primary appliance.

Suspending an HA appliance

If an HA appliance is suspended using the **sethagrp -s** the following sequence of notifications is seen:

```
AMQ3911W: HA group suspended on the local appliance  
AMQ3910W: HA group suspended on remote appliance 'Appliance1'
```

Both appliances report which appliance is being suspended.

```
AMQ3569I: HA role for queue manager QM_HA is none
```

Any queue managers running in the primary role on the suspended appliance are stopped and now have no HA role. Any queue managers in the secondary role are also reported as having no role.

```
AMQ3595W: HA status for queue manager QM_HA is 'Inactive'
```

For a brief period, any queue manager that is switching over reports Inactive, showing that both appliances are in the secondary role.

```
AMQ3902I: HA role for queue manager QM_HA is primary
```

The relocated queue manager is then promoted to Primary and started.

```
AMQ3576E: HA replication to remote appliance 'Appliance1' for queue manager 'QM_HA' using  
interface 'eth21' is unavailable
```

This notification reports that HA replication is no longer available as a consequence of the appliance suspending and stopping the data replication resources for all HA queue managers.

```
AMQ3591W: HA status for queue manager QM_HA is 'This appliance in standby mode'  
AMQ3590W: HA status for queue manager QM_HA is 'Secondary appliance in standby mode'
```

As the appliance that is suspending completes the job of stopping all the HA and data replication resources, both appliances report that the HA queue manager is in standby mode.

Resuming an HA appliance

When a suspended appliance is resumed using the **sethagr** **-r** command, the following sequence of notifications is seen:

```
AMQ3909I: HA group resumed on the local appliance  
AMQ3908I: HA group resumed on remote appliance 'Appliance1'
```

Both appliances report that the suspended node is being resumed.

```
AMQ3577I: HA replication to remote appliance 'Appliance1' for queue manager 'QM_HA' using  
interface 'eth21' is available
```

As the appliance restarts, the HA and data replication resources for each queue manager can again begin replication.

```
AMQ3592E: HA status for queue manager QM_HA is 'Inconsistent'  
AMQ3598W: HA status for queue manager QM_HA is 'Synchronization in progress'  
AMQ3599I: HA status for queue manager QM_HA is 'Normal'
```

As the data replication reconnects, the resuming appliance recognizes that the queue manager data is not up to date with the remote appliance and begins synchronization.

```
AMQ3903I: HA role for queue manager QM_HA is secondary  
AMQ3595W: HA status for queue manager QM_HA is 'Inactive'  
AMQ3902I: HA role for queue manager QM_HA is primary
```

Any queue managers that have their preferred location on the resumed appliance are stopped and demoted to secondary to allow them to be promoted to Primary and started on their preferred location.

Losing communication with remote appliance

There are several different notifications that can be seen when communication is lost between HA appliances, depending on which network interfaces are affected.

Replication interface

```
AMQ3576E: HA replication to remote appliance 'Appliance1' for queue manager 'QM_HA' using  
interface 'eth21' is unavailable  
AMQ3594E: HA status for queue manager QM_HA is 'Remote appliance unavailable'
```

Losing communication for replication is reported and results in a state change to Remote appliance unavailable. From this point the data on the secondary appliance becomes increasingly out of date.

Heartbeat interfaces

```
AMQ3904E: HA heartbeat connection to secondary appliance 'Appliance1' using interface  
'eth13' is unavailable
```

The appliance notifies when communication is unavailable for a single heartbeat link, that is, Primary interface (eth13) or Alternate interface (eth17). The appliance HA configuration can tolerate loss of a single heartbeat link.

```
AMQ3906S: HA secondary appliance 'Appliance1' is unavailable
AMQ3597E: HA status for queue manager QM_HA is 'Secondary appliance unavailable'
```

If both links are unavailable, additional notification is made of the remote appliance being unavailable, which is also reported for each affected HA queue manager.

```
AMQ3902I: HA role for queue manager QM_HA is primary
```

If the remote appliance is unavailable, all queue managers currently in the secondary role are promoted to primary and started.

Regaining communication with remote appliance

As with losing communication interfaces with the remote appliance, there are different notifications as each communications interface between HA appliances becomes available.

Regaining a heartbeat interface

```
AMQ3905I: HA heartbeat connection to secondary appliance 'Appliance1' using interface
'eth13' is available
```

The appliance notifies when each heartbeat connection, that is, Primary interface or Alternate interface, becoming available.

```
AMQ3907I: HA secondary appliance 'Appliance1' is available
```

If regaining a heartbeat connection allows HA coordination between the appliances to resume, the appliance notifies that the secondary appliance is available.

Regaining Replication

```
AMQ3577I: HA replication to remote appliance 'Appliance1' for queue manager 'QM_HA' using
interface 'eth21' is available
```

A notification is produced for each queue manager when the data replication connection is re-established.

Typically, this will be shortly followed by the secondary appliance reporting the queue manager data as inconsistent then synchronizing from the primary.

Partitioned behavior

If all HA heartbeat and replication links are lost but the remote appliance is still operating, the same sequence of states occurs on both appliances, resulting in the queue manager running simultaneously on both. This situation results in the queue manager data deviating in a manner that cannot be reconciled.

```
AMQ3596S: HA status for queue manager QM_HA is 'Partitioned'
```

When the replication links recover, the data replication resources detect the deviation in queue manager data and report that the queue manager has become Partitioned. If heartbeat links are available, this will be reported on both appliances. Duplicate notifications are possible if both appliances simultaneously detect the deviation.

On detecting the partitioned state, the replication connection becomes unavailable.

Additionally, one of the running instances is stopped and demoted to the secondary role. Typically, the queue manager remains running on the preferred node.

Resolving a partitioned state

After following the steps in “Resolving a partitioned problem in a high availability configuration” on page 561 to identify which appliance holds the 'winning' data for partition resolution, use the **makehaprimary** command on the identified appliance.

```
AMQ3903I: HA role for queue manager QM_HA is secondary
AMQ3902I: HA role for queue manager QM_HA is primary
```

If the **makehaprimary** command is run on the current secondary appliance, the queue manager instance is demoted on the primary appliance to allow it to be promoted, and started, where the command has been issued.

```
AMQ3577I: HA replication to remote appliance 'Appliance1' for queue manager 'QM_HA' using
interface 'eth21' is available
AMQ3592E: HA status for queue manager QM_HA is 'Inconsistent'
AMQ3598W: HA status for queue manager QM_HA is 'Synchronization in progress'
AMQ3599I: HA status for queue manager QM_HA is 'Normal'
```

Data replication is then reestablished between the appliances, with the data to be replaced marked as Inconsistent, allowing synchronization from the chosen primary node.

```
AMQ3595W: HA status for queue manager QM_HA is 'Inactive'
AMQ3903I: HA role for queue manager QM_HA is secondary
AMQ3902I: HA role for queue manager QM_HA is primary
```

After synchronization has completed, if **makehaprimary** was not issued on the preferred node, the queue manager is moved back to the preferred node.

DR notification examples

This topic illustrates common sequences of notifications that might be issued by a Disaster Recovery (DR) configuration.

Note: When a high availability (HA) queue manager is configured with DR, the current HA secondary appliance will notify only of changes to the DR role (other DR messages do not appear on the secondary appliance).

Configuring primary DR queue manager

When an appliance queue manager is enabled for disaster recovery, the following sequence of notifications are produced to indicate progress.

The following notifications are seen on the appliance where the **crtdrprimary** command is issued:

Replication notifications

```
AMQ3580E: DR status for queue manager QM_DR is 'Inconsistent'
```

As the data replication resources are created they are marked as Inconsistent to indicate the data is yet to be replicated to the remote appliance.

```
AMQ3581W: DR status for queue manager QM_DR is 'Remote appliance not configured'
```

When the DR primary instance is first created the remote appliance is reported as not configured, this will remain so until the generated **crtdrsecondary** command is issued on the remote appliance.

Disaster recovery (DR) notifications

```
AMQ3900I: DR role for queue manager QM_DR is primary
```

The appliance on which the **crtdrprimary** command is issued is promoted to the DR primary role.

Configuring secondary DR queue manager

When the **crtdrsecondary** command is issued on the remote appliance the following sequence of notifications is seen:

Disaster recovery (DR) notifications

```
AMQ3901I: DR role for queue manager QM_DR is secondary
```

The remote appliance takes on the DR secondary role.

Replication notifications

```
AMQ3580E: DR status for queue manager QM_DR is 'Inconsistent'
```

The data replication resource is marked as inconsistent to indicate that the DR secondary appliance needs to synchronize the queue manager data from the DR primary appliance.

```
AMQ3585W: DR status for queue manager QM_DR is 'Synchronization in progress'
```

After the data replication resources on both appliances connect, they begin synchronizing the data from the primary appliance to the secondary appliance. If replication is interrupted, for example, by interruption of the network, the state returns to Inconsistent.

```
AMQ3586I: DR status for queue manager QM_DR is 'Normal'
```

After the secondary appliance has synchronized, all disk writes from the primary appliance are asynchronously replicated to the secondary appliance. This is the Normal state.

Losing communication with remote appliance

When communication is lost between DR appliances the following notifications are seen:

```
AMQ3574E: DR replication to remote appliance 'Appliance2' for queue manager 'QM_DR' using  
interface 'eth20' is unavailable  
AMQ3582E: DR status for queue manager QM_DR is 'Remote appliance unavailable'
```

Losing communication for replication is reported for each affected queue manager and results in a state change to Remote appliance unavailable. From this point the data on the secondary appliance becomes increasingly out of date.

Regaining communication with remote appliance

As with losing data replication communication with the remote appliance, the following notification is seen when communication with the remote appliance becomes available:

```
AMQ3575I: DR replication to remote appliance 'Appliance2' for queue manager 'QM_DR' using  
interface 'eth20' is available
```

A notification is produced for each queue manager when the data replication connection is re-established.

Typically, this is shortly followed by the secondary appliance reporting that the queue manager data as Inconsistent then synchronizing from the primary.

Note:

When a high availability (HA) queue manager is configured with DR, if the DR recovery appliance issues notification AMQ3574E or AMQ3575I it will always give the name of the HA appliance on which the **crtdrprimary** command was issued as the remote appliance.

Switching over to a recovery appliance

Following the loss of the primary queue manager at the main site, you might need to start the secondary queue manager at the recovery site.

```
AMQ3900I: DR role for queue manager QM_DR is primary
```

The queue manager is promoted to the Primary role.

If the recovery site is inconsistent when the **makedrprimary** command is issued (that is, synchronization failed from the main site), the queue manager reverts to the previous saved snapshot of its data before it can be promoted to primary.

```
AMQ3579W: DR status for queue manager QM_DR is 'Reverting to snapshot'  
AMQ3582E: DR status for queue manager QM_DR is 'Remote appliance unavailable'
```

After reversion to the previous snapshot, the status is updated to Remote appliance unavailable.

The queue manager data has now deviated from that on the main site appliance and is reported as Partitioned when the replication connection becomes available.

Partitioned behavior

If the replication link is lost but the remote appliance is still operating, the same sequence of states occurs on both appliances. It is then possible to issue the **makedrprimary** command on the secondary appliance, resulting in the queue manager running simultaneously on both. This situation results in the queue manager deviating in a manner that cannot be reconciled.

```
AMQ3584S: DR status for queue manager QM_DR is 'Partitioned'
```

When the replication link recovers, the data replication resources detect the deviation in queue manager data and report that the queue manager has become Partitioned. The partitioned state might be detected on only one appliance, or both.

On detecting the partitioned state, the replication connection becomes unavailable.

Resolving a partitioned state

After following the steps in [“Resolving a partitioned problem in a disaster recovery configuration”](#) on page 564 to identify which appliance holds the ‘winning’ data for partition resolution, use the **makedrsecondary** command on the ‘losing’ appliance.

```
AMQ3901I: DR role for queue manager QM_DR is secondary  
AMQ3582E: DR status for queue manager QM_DR is 'Remote appliance unavailable'
```

The losing appliance is demoted to the DR secondary role and the status updated to indicate Remote appliance unavailable.

To complete the partition resolution, issue the **makedrprimary** command on the ‘winning’ appliance.

```
AMQ3575I: DR replication to remote appliance 'Appliance2' for queue manager 'QM_DR' using  
interface 'eth20' is available
```

Data replication is then re-established between the appliances, with the data to be replaced marked as Inconsistent, allowing synchronization from the chosen primary node.

```
AMQ3580E: DR status for queue manager QM_DR is 'Inconsistent'  
AMQ3585W: DR status for queue manager QM_DR is 'Synchronization in progress'  
AMQ3586I: DR status for queue manager QM_DR is 'Normal'
```

After synchronization has completed, if the **makedrprimary** command was not issued on the preferred primary node, the queue manager can be moved back using the **makedrsecondary** and **makedrprimary** commands.

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 V9 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/AppName/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/AppName/amqs*
```

Related concepts

[“System topics for monitoring and activity trace” on page 552](#)

System topics in queue manager topic trees are used for resource monitoring and for application activity trace.

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 `mqt.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.

```
amqsact -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 417 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 ActivityCount 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.

Related concepts

[“Monitoring system resource usage” on page 522](#)

You can monitor the use of system resources by IBM MQ.

[“Application activity trace” on page 547](#)

Application activity trace produces detailed information about the behavior of applications that are connected to a queue manager.

Monitoring the appliance by using the REST management interface

You can use the REST management interface to monitor the status of the IBM MQ Appliance.

When you use the REST management interface for this purpose, you send HTTP requests to the REST interface port and receive JSON-formatted responses with a payload and indication of success or failure. You can incorporate requests into programs and so automate interaction with the appliance.

You must be a local user to use the REST management interface. If you have configured role based management to use LDAP or XML file user authentication, then configure a fallback user to access the REST interface (see [“Role based management” on page 439](#)).

The appliance has a number of 'status providers'. You can retrieve complete status provider data for all existing status provider classes and retrieve individual property values of each status provider. The following topic provides an example of retrieving status by using the REST interface. For a reference guide to the REST management interface, see [“REST management interface” on page 1000](#).

Example of retrieving status by using REST

There are a number of major steps involved in retrieving status from the IBM MQ Appliance by using the REST management interface.

Identify a required status class

To begin retrieving the required status provider data from the appliance, first identify the specific status provider class that you need. To identify the required status class name, access the REST management interface root URI by using a GET request to identify the status root URI:

```
GET https://mqhost.com:5554/mgmt/
```

You receive the following response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/"
    }
  },
  "config": {
    "href": "/mgmt/config/"
  },
  "domains": {
    "href": "/mgmt/domains/config/"
  },
  "status": {
    "href": "/mgmt/status/"
  }
}
```

```

    },
    "actionqueue": {
      "href": "/mgmt/actionqueue/"
    },
    "filestore": {
      "href": "/mgmt/filestore/"
    },
    "metadata": {
      "href": "/mgmt/metadata/"
    },
    "types": {
      "href": "/mgmt/types/"
    }
  }
}

```

You can identify the status root URI in the received response as `/mgmt/status/`. Then, to retrieve a list of all available status provider classes on the appliance, make the following request:

```
GET https://mqhost.com:5554/mgmt/status/
```

To identify the exact formatting of the status provider class name, you search the received response payload. The following listing shows some fragments of the received response:

```

{
  "_links": {
    "self": {
      "href": "/mgmt/status/"
    },
    "ActiveUsers": {
      "href": "/mgmt/status/{domain}/ActiveUsers"
    },
    "Battery": {
      "href": "/mgmt/status/{domain}/Battery"
    },
    "ConnectionsAccepted": {
      "href": "/mgmt/status/{domain}/ConnectionsAccepted"
    },
    ...
    "LogTargetStatus": {
      "href": "/mgmt/status/{domain}/LogTargetStatus"
    },
    "MQSystemResources": {
      "href": "/mgmt/status/{domain}/MQSystemResources"
    },
    "NDCacheStatus2": {
      "href": "/mgmt/status/{domain}/NDCacheStatus2"
    },
    ...
  }
}

```

Retrieve complete status data

After you identify the required status provider class name, you can retrieve the associated status data. To retrieve the data, you construct a URI of the form `/mgmt/status/domain/class_name`, replacing *domain* with the string "default" and *class_name* with the desired status provider class. The following request shows a URI to retrieve information from the `MQSystemResources` status provider within the default domain:

```
https://mqhost.com:5554/mgmt/status/default/MQSystemResources
```

The status provider returns the following information:

```

{
  "_links" : {
    "self" : {
      "href" : "/mgmt/status/default/MQSystemResources"
    },
    "doc" : {
      "href" : "/mgmt/docs/status/MQSystemResources"
    }
  },
  "MQSystemResources" : {

```

```

        "TotalStorage" : 15667,
        "UsedStorage" : 9216,
        "TotalErrorsStorage" : 1024,
        "UsedErrorsStorage" : 40,
        "TotalTraceStorage" : 2048,
        "UsedTraceStorage" : 281,
        "HAStatus" : "",
        "HAPartner" : ""
    }
}

```

The following request shows a URI to retrieve information from the DateTimeStatus status provider within the default domain:

```
https://mqhost.com:5554/mgmt/status/default/DateTimeStatus
```

The status provider returns the following information:

```

{
  "_links" : {
    "self" : {
      "href" : "/mgmt/status/default/DateTimeStatus"
    },
    "doc" : {
      "href" : "/mgmt/docs/status/DateTimeStatus"}
  },
  "DateTimeStatus" : {
    "time" : "Mon Oct 31 14:29:37 2016",
    "timezone" : "GMT",
    "tzspec" : "GMT0BST,M3.5.0/1:00,M10.5.0/2:00",
    "uptime2" : "3 days 03:19:14",
    "bootuptime2" : "3 days 03:21:39"
  }
}

```

Retrieve partial status data

You can also retrieve the value of a specific status provider property, instead of retrieving the status provider output in its entirety. To retrieve the value, you construct a URI of the form `/mgmt/status/domain/class_name/property_name`. You replace *domain* with the string "default", *class_name* with the required status provider class, and *property_name* with the specific property name as it appears in the complete status provider response. For example, you could enter the following URI to retrieve just the up time from the datetime status provider:

```
https://mqhost.com:5554/mgmt/status/default/DateTimeStatus/uptime2
```

The status provider returns the following information:

```

{
  "_links" : {
    "self" : {
      "href" : "/mgmt/status/default/DateTimeStatus/uptime2"
    },
    "doc" : {
      "href" : "/mgmt/docs/status/DateTimeStatus/uptime2"}
  },
  "DateTimeStatus" : {
    "uptime2" : "3 days 03:19:14",
  }
}

```

Related concepts

[“REST management interface” on page 1000](#)

The REST management interface provides access to the actions and to the configuration and status resources on the appliance.

Monitoring the IBM MQ Appliance by using SNMP

You can configure SNMP to monitor the appliance. The appliance supports SNMP versions 1, 2c, and 3.

When you configure SNMP on the appliance, you enable one or more SNMP managers to interrogate the appliance to retrieve information about its current state. The appliance objects that can be interrogated are defined in three MIB files. You can view the MIB files by using the web UI (see [“Viewing MIBs by using the web UI”](#) on page 229).

The appliance can also respond to events by generating traps (v1 and v2c) or notifications (v3). These traps or notifications can be sent to SNMP managers to inform them that the event has occurred.

You can configure SNMP on the appliance either by using the web UI or the command line interface. See [“SNMP Settings”](#) on page 227.

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 view the log files generated by IBM MQ and by the appliance itself. You can configure log targets so that certain types of logging information is written to specified destinations, such as a syslog system. (See [“IBM MQ appliance logs”](#) on page 497.)
- You can view error reports generated by the appliance. You can configure the appliance so that it automatically uploads a First Failure Data Capture (FFDC) to an FTP server or NFS mount. (See [“Appliance error reports”](#) on page 557)
- 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.

Related reference

[“mqrc”](#) on page 600

Display information about return codes.

[“runmqras”](#) on page 603

Gather diagnostic information together into a single archive to submit to IBM Support.

[“dspmqrte \(display route information\)”](#) on page 586

Determine the route that a message has taken through a queue manager network.

Appliance error reports

An appliance error report file contains the current configuration of the appliance and the current contents of the appliance log files.

You can generate an error report whenever you want, or you can configure a failure notification to be generated automatically when a failure occurs. The contents of error reports and where these reports are sent depend on the failure notification configuration.

V 9.1.5 By default, error reports are written to the `temporary:` directory (which is on RAM disk and discarded when the appliance shuts down or restarts). You can choose the upload option to specify that they are written to the `mqdiag:` directory (which is on RAID disk and so persistent). If you choose the upload option, you can specify more options as to what is included in the error report. This enables First Failure Data Capture (FFDC).

If you do not choose the upload option, you can choose the email option and simply specify an email recipient for an error report generated when the appliances shuts down or starts up. Note that this error report does not contain FFDC information.

To resolve a problem, IBM Support requests an error report.

Configuring failure notification using the web UI

With failure notification enabled, you can send or upload an error report after the IBM MQ Appliance returns to service from an outage. The error report can contain diagnostic details.

About this task

When you choose to upload error reports, you can configure First Failure Data Capture (FFDC) settings. An FFDC event triggers the error report and contains additional information. FFDC events include crashes, out-of-memory crashes, watchdogs, and other user-requested events.

When you choose to send email notifications, FFDC information is not included.

Procedure

1. Start the IBM MQ Appliance web UI, and click the administration icon .
2. Select **Device > Failure Notification**.
3. Ensure that **Enable administrative state** is selected.
4. Optionally enter a description in the **Comments** field.
5. Use the **Include Internal State** property to control whether error reports include a snapshot of the system.
6. Use the **Always on Startup** property to control whether to send an error report on each firmware restart.
7. Use the **Always on Shutdown** property to control whether to generate an error report on each firmware shutdown.
8. Select **Upload error report** to upload the error report to a specified destination. Further fields appear that enable you to specify the destination.
 - Use the **Background Packet Capture** property to control whether to execute an always-on background packet capture that captures network packets on all interfaces, which include the internal loopback interface. When an FFDC event is triggered, that error report includes these details. Although this feature is disabled by default, it is recommended that you enable it. If you enable this feature and manually request a packet capture, two packet-captures are running, which significantly affects performance.
 - Use the **Background Log Capture** property to control whether to execute an always-on background log message capture that captures all log and trace points in the code with minimal performance impact. Log messages that are captured are independent of the system logging configuration. When an FFDC event is triggered, that error report includes these details. Although this feature is disabled by default, it is recommended that you enable it.
 - Use the **Background Memory Trace** property to control whether to run memory-tracing when system memory usage becomes too low or too high. This feature enables background leak detection that captures all allocation call sites when the system detects a memory leak trend. When an FFDC event is triggered, that error report includes these details. Although this feature is disabled by default, it is recommended that you enable it.
 - In the **Report Destination Protocol** field, select the protocol to use to upload the error report.

To send as email:

 - a. Select **smtp**.
 - b. In the **SMTP Server** field, enter the host name or IP address of the remote SMTP server to which the error-report is sent. Click **Ping** to test the connection to the chosen server.
 - c. In the **Email Address** field, enter the email address of the recipient.
 - d. In the **Email Sender Address** field, optionally enter the email address of the sender. If not specified, the email address of the sender is the same as the recipient.

V 9.1.5 To write to the mqdiag: directory on the RAID disk, select **mqdiag** as the as the **Report Destination Protocol**.

To write the report to the temporary: directory, select **temporary** as the **Report Destination Protocol**. The report is written to temporary: by default, but if you select this option as the upload destination, you can select the extra information to include in the report.

9. When not uploading error reports, use the **Email Notification** property to control whether to send error reports to an email recipient when the appliance shuts down or starts up. This option does not include FFDC information in the report. When you select this option further fields enable you to specify details about the email destination:
 - In the **Location** field, enter the subject line of the email. Include a customer name or system identifier in this field.
 - In the **SMTP Server** field, enter the host name or IP address of the remote SMTP server to which the error-report is sent. Click **Ping** to test the connection to the chosen server.
 - In the **Email Address** field, enter the email address of the recipient.
 - In the **Email Sender Address** field, optionally enter the email address of the sender. If not specified, the email address of the sender is the same as the recipient.
10. In the **Report history** field, specify the number of error reports to maintain. When the number is reached, the next time an error report is written, it overwrites the oldest saved report.

Generating error reports

Generate an error report so that you can view the current configuration, status providers, and the current contents of various log files.

About this task

You can generate an error report from the Troubleshooting page of the IBM MQ Appliance web UI. IBM® Support can request that you to generate an error report to help diagnose a problem.

Procedure

1. Start the IBM MQ Appliance web UI, and click the administration icon .
2. Select **Debug > Troubleshooting**.
3. In the **Reporting** section, click **Generate Error Report**.

Viewing error reports

After generating an error report, you can view it.

About this task

You can only view the most recently generated report. If someone else generates an error report after you, your report is no longer available in this way.

If your error report is not the most recent report, you can download and extract older reports. Older reports are available as compressed files. The storage location is set during the configuration for failure notification.

Procedure

1. Start the IBM MQ Appliance web UI, and click the administration icon .
2. Select **Debug > Troubleshooting**.
3. In the **Reporting** section, click **View Error Report**.
4. Follow the prompts.

Sending error reports

How to send an error report to an email address.

About this task

The send error report tool sends the error report as a compressed file attachment to an email address. To send the report, the appliance must have a connection to an SMTP server that accepts the message.

When failure notification uses an SMTP server, do not use this task unless you want to send the error report to a different email recipient.

Procedure

1. Start the IBM MQ Appliance web UI, and click the administration icon .
2. Select **Debug > Troubleshooting**.
3. In the **Reporting** section, in the **Send Error Report** subsection, supply the following information.
 - In the **SMTP Server** field, enter the IP address or host name of the SMTP server.
 - In the **Subject** field, enter the subject line of the email.
 - In the **Email Address** field, enter the email address of the recipient.
4. Click **Send Error Report**.

Using trace

You can use the **strmqtrc** and **endmqtrc** commands on the command line to start and end tracing. You can use the **dspmqrtrc** command to format and view the trace.

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 \(Start trace\)” on page 613](#).

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 \(end trace\)” on page 599](#).

After you generate the trace files, you can download them from the appliance. **V 9.1.5** You can also use the **dspmqrtrc** command to format and view them on 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. **V 9.1.5** Alternatively, you can use the **dspmqrtrc** command to format and view the trace file on the appliance.

Related reference

[“strmqtrc \(Start trace\)” on page 613](#)

Start trace at a specified level of detail, or report the level of tracing in effect.

[“endmqtrc \(end trace\)” on page 599](#)

End trace for some or all of the entities that are being traced.

[“runmqras” on page 603](#)

Gather diagnostic information together into a single archive to submit to IBM Support.

[“copy” on page 788](#)

This command copies a file to or from the appliance.

[“dspmqtrc \(display formatted trace\)” on page 593](#)

Format and display IBM MQ trace.

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 settings icon  in the IBM MQ Console title bar
2. In the Settings window, click **Enable** for **Browser trace**.

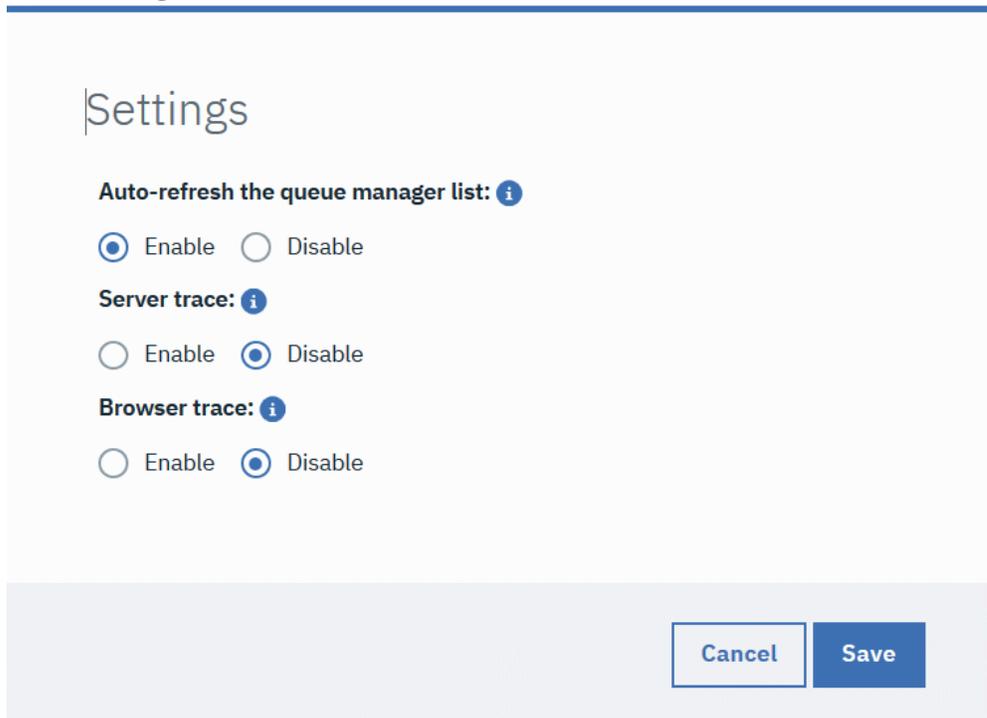


Figure 76. The settings window

3. Click **Save**.

You can then re-create the problem that you are trying to troubleshoot and capture the results in the trace. You then disable trace once more by selecting **Disable** for the **Browser trace** option.

Capture save the trace in a file by using the `runmqras -section trace` command from the command line (see [“Using trace” on page 560](#)). You can download the trace file from the File Management window of the IBM MQ Appliance web UI (see [“Managing files by using the IBM MQ Appliance web UI” on page 402](#)). The trace files is stored in the `mqdiag://` directory.  You can also format and view the trace by using the `dspmqtrc` command, see [“dspmqtrc \(display formatted trace\)” on page 593](#).

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 both the primary and secondary connections are lost, the queue manager will run on both appliances at the same time.

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). When the first connection is restored (either primary or secondary), the queue manager is stopped on one appliance but continues to run on the other. The HA status is shown as `partitioned`.

Choosing the 'winner'

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 one appliance, and the data replicated to the other appliance.

To help you decide, you can run the `status` command for the affected queue manager on each appliance. The status command returns an HA status of `partitioned` together with a report of how much out-of-sync data the appliance has for that queue manager. See [“Viewing the status of a high availability queue manager”](#) on page 278.

You can also view the actual data associated with the affected queue manager on each appliance by interrogating the state of each queue manager. Stop all your HA queue managers and disable HA control by issuing the following command for each HA queue manager:

```
endmqm -w qmgr
```

Then follow these steps to run each affected queue manager on each appliance, outside of HA control, so that you can interrogate the state of each version of the queue manager:

1. Display the HA status to confirm that the `HA control` field shows disabled, and that the queue manager is ended on both appliances:

```
status qmgr
```

See [“Viewing the status of a high availability queue manager”](#) on page 278. (If the queue manager is running on the other appliance, end it there too.)

2. On the appliance that is the primary for the queue manager, restart the queue manager outside of HA control:

```
strmqm -ns qmgr
```

This command starts the queue manager without starting the listener, to prevent any applications connecting. (If you need a listener in order to connect to the queue manager from, for example, IBM MQ Explorer, then you should start a listener manually by using a `RUNMQSC` command. Start the listener on a different port to the normal one so that applications cannot connect.)

3. Display the HA status again and verify that the queue manager is running and that `HA control` still shows disabled:

```
status qmgr
```

4. Interrogate the state of the queue manager, for example, by browsing messages.
5. End the queue manager once more:

```
endmqm -w qmgr
```

6. Display the HA status again and verify that the queue manager is ended on both appliances:

```
status qmgr
```

7. Suspend the HA group on this appliance by entering the following command:

```
sethgrp -s
```

8. Display the HA status of this appliance to confirm that it is in the standby state:

```
dsphagip
```

9. Log in to the other appliance and start the queue manager:

```
strmqm -ns qmgr
```

(This command starts the queue manager without the listener. If you need a listener, see the advice in step 2.)

10. Check the status and confirm that the queue manager is running and that HA control is still disabled:

```
status qmgr
```

11. Interrogate the state of the queue manager on that appliance, for example, by browsing messages.

12. End the queue manager:

```
endmqm -w qmgr
```

13. Display the HA status and confirm that HA control is still disabled and that the queue manager is ended on both appliances.

14. Resume the suspended appliance:

```
sethagip -r
```

15. Display the HA status and confirm the suspended appliance is now active, and HA control is still disabled, and HA status is `Partitioned` on both appliances.

```
dsphagip  
status qmgr
```

Implementing your choice of winning appliance

You identify the winner by running the following commands on the chosen appliance:

1. Make the queue manager on the winning appliance the primary:

```
makehprimary HAQMName
```

Where *HAQMName* is the name of the queue manager. The data is then synchronized between the queue managers on the winning appliance and the other appliance.

2. Check the HA status to confirm that synchronization has finished and a status of `Normal` is reported on both appliances:

```
status qmgr
```

3. Start the queue manager with HA control:

```
strmqm qmgr
```

Whichever appliance that you start the queue manager on, it should run on the appliance that has been designated as its preferred appliance.

4. Display the HA status on both appliances and confirm HA control is enabled and the queue manager is active on only the preferred appliance:

```
status qmgr
```

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 ... > **High Availability (HA) > Resolve partitioned data.**
3. Confirm that you want the current appliance to become the primary data source for the queue manager.

Other situations

If the two appliances lose the replication interface, the HA status is reported as Remote appliance(s) unavailable. 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, 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 382.

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):

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 required 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
```

Related concepts

[“Switching back to the main appliance” on page 382](#)

When the disaster has been resolved, and the main appliance restored, you can revert to running the queue manager on your main appliance.

Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

[“makedrsecondary” on page 685](#)

Prevents a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.

Resolving an HA queue manager left in an indeterminate state

If a power failure occurs when adding a queue manager to a high availability (HA) group, it can leave the queue manager in an indeterminate state.

HA commands can take some time to run. A power failure, or similar disruption, occurring when you are adding an existing queue manager to an HA group can leave the queue manager in an indeterminate state where it is running on neither appliance in the HA pair.

You can resolve this situation by using one of the following methods (the 'local' appliance is the one where you issued the commands that were interrupted).

First, recover the HA status of the queue manager from both appliances by using the **status QMName** command.

- If the queue manager on the local appliance is in a non-HA state and the remote appliance is in a HA state, then complete the following steps:

1. On the remote appliance, enter the following command:

```
dltmqm QMName
```

2. On the local appliance, enter the following command:

```
sethgrp -i QMName
```

The queue manager will be added to the HA group and run on the local appliance.

- If the queue manager on the local appliance is in an indeterminate HA state and the remote appliance is in an HA state, then complete the following steps:

1. On the remote appliance, enter the following command:

```
dltmqm QMName
```

You might need to repeat this command several times.

2. On the local appliance, enter the following command:

```
sethgrp -e QMName
```

The queue manager will run on the local appliance as a stand-alone queue manager.

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.

URI	Permissions	Description
mqbackup://	copy from copy to	Used for: <ul style="list-style-type: none"> – user backup and restore – certificate backup and restore – dmpmqcfg output – CCDTs – V 9.1.4 automatic configuration of members of uniform cluster – V 9.1.4 runmqsc input files
mqdiag://	copy from	Used for: <ul style="list-style-type: none"> – runmqras output – amqrfdm – amqspdbg

URI	Permissions	Description
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 security management commands” on page 629)
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. ▶ V 9.1.4 The directory can also be used for runmqsc input files.
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`.

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.

▶ M 2000 M 2001 Troubleshooting SAN problems

What happens when you lose SAN connectivity.

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances.

If you have configured a queue manager to use SAN storage (rather than local appliance storage), you will encounter problems if you lose connection to the SAN fabric.

In all cases it is recommended that you use multipath connections to SAN, and enable multipath in your SAN configurations on the appliance. This reduces the likelihood of losing access to the remote storage.

If you do lose connectivity across all paths, you might see errors in your application (for example, reporting object damaged), and first failure data captures (FDCs) might be recorded on the appliance. Additionally, the Fibre-Channel-Volume object transitions to Op-State Down (you can view the operational state of the object by using the **show fibre-channel-volume-status** command). The state transition causes the queue manager that is using that Fibre-Channel-Volume to transition from a status such as Running or Ending to Status Not Available.

After the routes to the SAN storage are restored, the Fibre-Channel-Volume must be disabled then re-enabled to transition from Op-State Down to Op-State Up. This transition causes the queue manager status to change to Ended Unexpectedly. You should then be able to start the queue manager to resume normal operations.

If you can restore SAN connectivity but the primary appliance that was previously running the queue manager on SAN cannot be recovered, you can recreate the queue manager on a different appliance. The SAN Volume containing the queue manager data must be made available to the other appliance.

You should create the Fibre-Channel-Volume using the same attributes as the original volume. Create the volume in the disabled state, because the newly-created volume uses the same LUID as the original, which will still have the primary appliance's locks on that volume. You must issue the `fibre-channel-unlock-volume volume_name` to clear the primary appliance's locks.

Having cleared the locks, the next step is to enable the volume for use. To recover the queue manager you must run the `addmqm -fc volume_name -m queue_manager_name` command from the `mqcli` prompt. Running this command restores the queue manager on the new appliance. The queue manager has the ended state.

Unlocking volumes

You might require to unlock a volume, for example, if it has been left locked by an appliance that stopped abruptly while having the volume enabled. You can unlock the volume from another appliance to take over the work from the failed appliance.

You unlock a volume by using the **fibre-channel-unlock-volume** command (see [“fibre-channel-unlock-volume”](#) on page 797).

The appliance that unlocks a volume must be zoned such that it can see the volume.

When a volume is locked, any other appliance should be able to clear the locks whether the volume was defined as multipath or non-multipath. If, however, the locks are cleared on a non-multipath volume by the appliance on which the volume became locked, the volume must remain defined as non-multipath. Otherwise a registration conflict occurs.

Related concepts

[“Configuring SAN storage”](#) on page 300

You can configure the IBM MQ Appliance so that queue managers can use SAN storage.

Related reference

[“addmqm \(add existing queue manager\)”](#) on page 644

Add an existing queue manager that uses SAN storage.

[“fibre-channel-unlock-volume”](#) on page 797

This command unlocks the specified volume.

[“show fibre-channel-volume-status”](#) on page 890

This command displays the current status of the fibre channel volumes.

Problems resizing queue managers

You might encounter a problem where you have used the **setmqsize** command to increase the size of a queue manager's file system, but the size has not actually increased.

The **setmqsize** command uses a two-stage operation when increasing the file system size. If the operation is interrupted before both stages have completed (for example, by a power failure) your queue manager can be left in the situation where it reports having a file system of the new size, but actually still has the old file system size.

To remedy this situation, re-run the **setmqsize** command, specifying the desired file system size.

Related reference

[“setmqsize \(increase queue manager file system size\)” on page 659](#)
Increase the size of the file system allocated to a queue manager.

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 603](#) 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_HHmss.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_HHmss.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 **runmqras -section all** with caution, consider using the `qmlist` parameter to limit the amount of output generated. See [“runmqras” on page 603](#) for details.
- Use **runmqras -section trace** to collect the trace files.
- Use **runmqras -section format** to generate and include formatted 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.

Related reference

[“runmqras” on page 603](#)

Gather diagnostic information together into a single archive to submit to IBM Support.

▶ M2001 Recovering from hardware failures, M2001 model

In the situation where you experience a hardware failure on an appliance, there are various steps you can take to get your queue managers running again as soon as possible.

If you detect a hardware failure, for example, indicated by the **test hardware** command (see [“test hardware command” on page 96](#)), contact IBM support for assistance. However, in some cases and depending on availability of spare components, there may be local actions you can take to quickly restore your appliance queue managers to operation.

The following topics describe scenarios that might occur and give step by step instructions to recover from them.

Related concepts

[“Diagnosing your appliance” on page 92](#)

How to diagnose problems in your appliance.

[“Troubleshooting your appliance” on page 99](#)

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.

Appliance fails, both disks unaffected

In this scenario, your appliance has experienced a failure that stops it operating, such as main board or RAID board failure, but the disks themselves are unaffected.

If you have a spare appliance, you can swap your two disks into the spare, and restart running your queue managers with the minimum of disruption.

To swap your disks to a new appliance:

1. Shut down both appliances (assuming both are running) by using the following command:

```
mqa# shutdown poweroff
```

2. Remove the disks from the first appliance, see [“Replacing a solid state disk drive module - M2001 appliances” on page 110](#) for instructions on how to remove disks.
3. Fit the disks into the second appliance. You might wish to place the disks in the equivalent slots to the appliance that they were removed from, but as the contents of each disk is identical, positioning is not significant.
4. Power up the second appliance.
5. The RAID controller detects the new disks and places them online automatically. You can check for disk status by entering the following commands:

```
mqa# config  
mqa (config)# show raid-physical-drive
```

The `state` field shows the state `online`.

6. If the state of either disk is `unconfiguredGoodForeign`, you can remedy this situation by entering the following command:

```
mqa (config)# raid-make-hot-spare raid0
```

The **raid-make-hot-spare** command discards any state (data) on the `unconfiguredGoodForeign` disk and activates it as part of the local RAID volume. You can monitor the progress of the resulting rebuild by using the **show raid-physical-drive** command again.

If any of the disks have the state `unconfiguredBad` or `unconfiguredBadForeign`, these states cannot be recovered by you and you must refer to IBM Support for help recovering the data.

7. Enter the following commands to check the state of your queue managers:

```
mqa# mqcli  
mqa(mqcli)# dspmq
```

If this procedure does not go as expected, contact IBM Support.

Appliance fails, one disk unaffected

In this scenario, your appliance has experienced a failure that stops it operating, such as main board or RAID board failure, one disk out of the pair is good.

If you have a spare appliance, you can insert your one good disk into the spare appliance and replicate its contents to another good disk to make a RAID pair. You can then restart and run your queue managers with the minimum of disruption.

This scenario assumes that the spare appliance has no disks of its own. First, you install the disks from which you want to recover the data, and check whether the disk fault is still detected on the new appliance. If a disk is still shown as having a fault, you replace it with a new disk and configure that.

To swap your disks to a new appliance:

1. Shut down the failed appliance by using the following command:

```
mqa# shutdown poweroff
```

2. Remove the both disks from the appliance. See [“Replacing a solid state disk drive module - M2001 appliances” on page 110](#) for instructions on how to remove disks.
3. Fit the disks into the second appliance. You might wish to place the disks in the equivalent slot to the appliance that they were removed from, but positioning is not significant.
4. Power up the second appliance.
5. The RAID controller detects the new disks automatically. Use the following command to check their status:

```
mqa# config  
mqa (config)# show raid-physical-drive
```

If the state of both disks shows `online`, then you need take no further action, you can proceed to step 10 to check your queue managers. If one disk still has a fault, however, proceed to the next step.

6. Shut down then power down the appliance, remove the faulty disk and fit a replacement disk in its place. Then power up the appliance again.
7. The RAID controller detects the new disk and begins a RAID rebuild automatically. The rebuild can take some hours to complete, you can check for completion by entering the **show raid-physical-drive** command again.
8. When the rebuild is complete for both disks, the state field should show the state `online`, and you can continue to step 10 to check the state of your queue managers. If the state of either disk is `unconfiguredGoodForeign`, you can remedy this situation by entering the following command:

```
mqa (config)# raid-make-hot-spare raid0
```

The **raid-make-hot-spare** command discards any state (data) on the `unconfiguredGoodForeign` disk and activates it as part of the local RAID volume. Again you can monitor the progress of the rebuild by using the **show raid-physical-drive** command.

If any of the disks have the state `unconfiguredBad` or `unconfiguredBadForeign`, these states cannot be recovered by you and you must refer to IBM Support for help recovering the data.

9. Enter the following commands to check the state of your queue managers:

```
mqa# mqcli  
mqa(mqcli)# dspmq
```

If this procedure does not go as expected, contact IBM Support.

Appliance operational, one disk in RAID pair fails

If one of your disks fails, you can replace it with another disk and resume operation.

Follow this procedure to replace a failed disk. You can use a good disk from another appliance or a new, replacement disk. The RAID system replicates the contents of the existing disk to the new disk.

1. Shut down the appliance by using the following command:

```
mqa# shutdown poweroff
```

If you are taking a disk from another appliance, you must shut down the donor appliance too.

2. Remove the failed disk from your first appliance. Follow the instructions in [“Replacing a solid state disk drive module - M2001 appliances”](#) on page 110.
3. If you are reusing a disk from another appliance, remove that disk.
4. Insert the good disk in the appliance to replace your failed disk.
5. Power up the appliance that you have replaced the disk in.
6. When a new disk with an unknown status is detected, a RAID rebuild begins automatically. This procedure can take some hours, you can check for completion by entering the following commands:

```
mqa# config  
mqa (config)# show raid-physical-drive
```

When complete, the state field contains the following text:

```
online
```

If a rebuild does not begin automatically, you can force one by entering the following command:

```
mqa (config)# raid-make-hot-spare raid0
```

The **raid-make-hot-spare** command discards any state (data) on the new disk and activates it as part of the local RAID volume.

7. Enter the following commands to check the state of your queue managers:

```
mqa(config)# exit  
mqa# mqcli  
mqa(mqcli)# dspmq
```

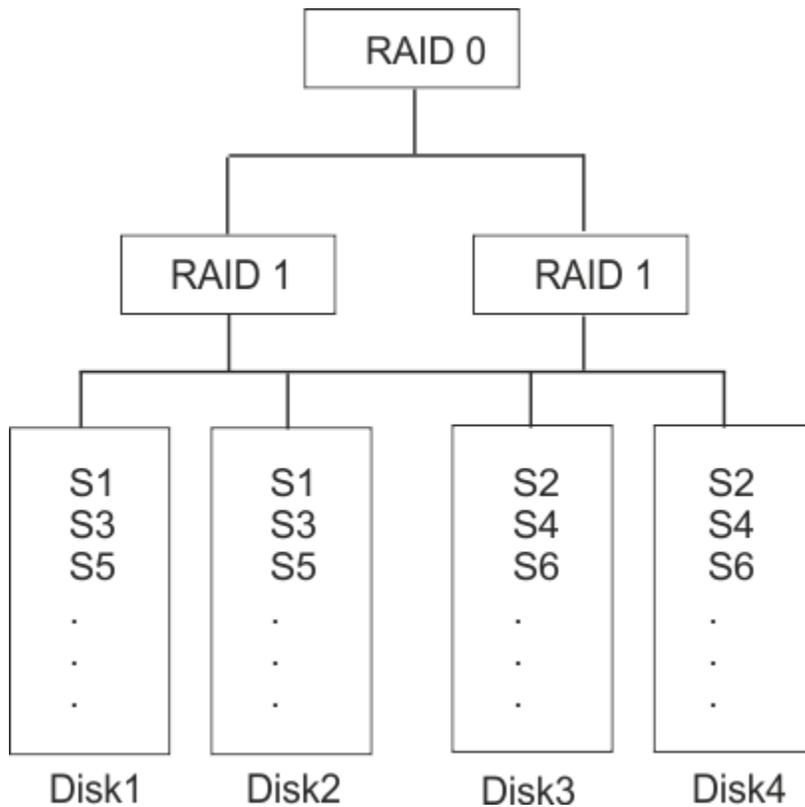
If this procedure does not go as expected, contact IBM Support.

M 2002 Recovering from hardware failures, M2002 model

In the situation where you experience a hardware failure on an appliance, there are various steps that you can take to get your queue managers running again as soon as possible.

If you detect a hardware failure, for example, indicated by the **test hardware** command (see [“Hardware testing from the CLI”](#) on page 148), contact IBM support for assistance. However, in some cases and depending on availability of spare components, there might be local actions that you can take to quickly restore your appliance queue managers to operation.

The M2002 appliance has four solid-state disks (SSDs), in a RAID10 configuration. Data is striped and mirrored across the four disks, as illustrated in the following diagram.



This configuration means that sometimes data is retrievable if two of the four SSDs fail, and sometimes not, depending on which disks fail. For example, in the configuration in the illustration, data could be retrieved if disks 1 and 3 fail, but not if disks 1 and 2 fail.

The following topics describe scenarios that might occur and give step-by-step instructions to recover from them.

Related concepts

[“Diagnostics and the appliance” on page 146](#)

The appliance and its firmware provide various ways to help you diagnose problem with hardware components.

Related tasks

[“Troubleshooting your appliance” on page 151](#)

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.

► M 2002 Appliance fails, all disks unaffected

In this scenario, your appliance has experienced a failure that stops it operating, such as main board or RAID board failure, but the disks themselves are unaffected.

If you have a spare appliance, you can swap your four disks into the spare, and restart running your queue managers with the minimum of disruption.

To swap your disks to a new appliance:

1. Shut down both appliances (assuming both are running) by using the following command:

```
mqa# shutdown poweroff
```

2. Remove the disks from the first appliance, see [“Replacing a solid-state disk drive module” on page 161](#) for instructions on how to remove disks.
3. Fit the disks into the second appliance. You must place the disks in the same position as the one they were removed from.

4. Power up the second appliance.
5. The RAID controller detects the new disks and places them online automatically. You can check for disk status by entering the following commands:

```
mqa# config
mqa (config)# show raid-physical-drive
```

The `state` field shows the state online.

6. If the state of any disk is `unconfiguredGoodForeign`, you can remedy this situation by entering the following command:

```
mqa (config)# raid-make-hot-spare raid0
```

The **raid-make-hot-spare** command discards any state (data) on the `unconfiguredGoodForeign` disk and activates it as part of the local RAID volume. You can monitor the progress of the resulting rebuild by using the `show raid-physical-drive` command again.

If any of the disks have the state `unconfiguredBad` or `unconfiguredBadForeign`, these states cannot be recovered by you and you must refer to IBM Support for help recovering the data.

7. Enter the following commands to check the state of your queue managers:

```
mqa# mqcli
mqa(mqcli)# dspmq
```

If this procedure does not go as expected, contact IBM Support.

M 2002 Appliance fails, three disks unaffected

In this scenario, your appliance has experienced a failure that stops it operating, such as main board or RAID board failure, three disks out of the four are good.

If you have a spare appliance, you can insert your good disks into the spare appliance and replicate the contents to another good disk to re-create the RAID10 configuration. You can then restart and run your queue managers with the minimum of disruption.

This scenario assumes that the spare appliance has no disks of its own. First, you install the disks from which you want to recover the data, and check whether the disk fault is still detected on the new appliance. If a disk is still shown as having a fault, you replace it with a new disk and configure that.

To swap your disks to a new appliance:

1. Shut down the failed appliance by using the following command:

```
mqa# shutdown poweroff
```

2. Remove the disks from the appliance. See [“Replacing a solid-state disk drive module” on page 161](#) for instructions on how to remove disks.
3. Fit the disks into the second appliance. You must place the disks in the equivalent slots to the appliance that they were removed from.
4. Power up the second appliance.
5. The RAID controller detects the new disks automatically. Use the following command to check their status:

```
mqa# config
mqa (config)# show raid-physical-drive
```

If the state of all four disks shows `online`, then you need take no further action, you can proceed to step 10 to check your queue managers. If one disk still has a fault, however, proceed to the next step.

6. Shut down then power down the appliance, remove the faulty disk and fit a replacement disk in its place. Then power up the appliance again.

7. The RAID controller detects the new disks and begins a RAID rebuild automatically. The rebuild can take some hours to complete, you can check for completion by entering the **show raid-physical-drive** command again.
8. When the rebuild is complete for all four disks, the state field should show the state online, and you can continue to step 10 to check the state of your queue managers. If the state of any disk is unconfiguredGoodForeign, you can remedy this situation by entering the following command:

```
mqa (config)# raid-make-hot-spare raid0
```

The **raid-make-hot-spare** command discards any state (data) on the unconfiguredGoodForeign disk and activates it as part of the local RAID volume. Again you can monitor the progress of the rebuild by using the **show raid-physical-drive** command.

If any of the disks have the state unconfiguredBad or unconfiguredBadForeign, these states cannot be recovered by you and you must refer to IBM Support for help recovering the data.

9. Enter the following commands to check the state of your queue managers:

```
mqa# mqcli  
mqa(mqcli)# dspmq
```

If this procedure does not go as expected, contact IBM Support.

Appliance operational, one disk in RAID set fails

If one of your disks fails, you can replace it with another disk and resume operation.

Follow this procedure to replace a failed disk. You can use a good disk from another appliance or a new, replacement disk. The RAID system replicates the contents of the existing disk to the new disk.

1. Shut down the appliance by using the following command:

```
mqa# shutdown poweroff
```

If you are taking a disk from another appliance, you must shut down the donor appliance too.

2. Remove the failed disk from your first appliance. Follow the instructions in [“Replacing a solid state disk drive module - M2001 appliances”](#) on page 110.
3. If you are reusing a disk from another appliance, remove that disk.
4. Insert the good disk in the appliance to replace your failed disk.
5. Power up the appliance that you have replaced the disk in.
6. When a new disk with an unknown status is detected, a RAID rebuild begins automatically. This procedure can take some hours, you can check for completion by entering the following commands:

```
mqa# config  
mqa (config)# show raid-physical-drive
```

When complete, the state field contains the following text:

```
online
```

If a rebuild does not begin automatically, you can force one by entering the following command:

```
mqa (config)# raid-make-hot-spare raid0
```

The **raid-make-hot-spare** command discards any state (data) on the new disk and activates it as part of the local RAID volume.

7. Enter the following commands to check the state of your queue managers:

```
mqa(config)# exit  
mqa# mqcli  
mqa(mqcli)# dspmq
```

If this procedure does not go as expected, contact IBM Support.

Reference

You can use the reference information in this section to accomplish the tasks that address your business needs.

Command reference

The command line is the accessible interface for the IBM MQ Appliance. The commands include methods to configure the appliance itself. The commands also include methods to manage IBM MQ objects.

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 207](#).

You must enter MQ administration mode before issuing commands. To enter MQ administration mode, type `mqcli` on the command line. You see the prompt `mqc(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 [Railroad diagrams](#).

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:

```
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`

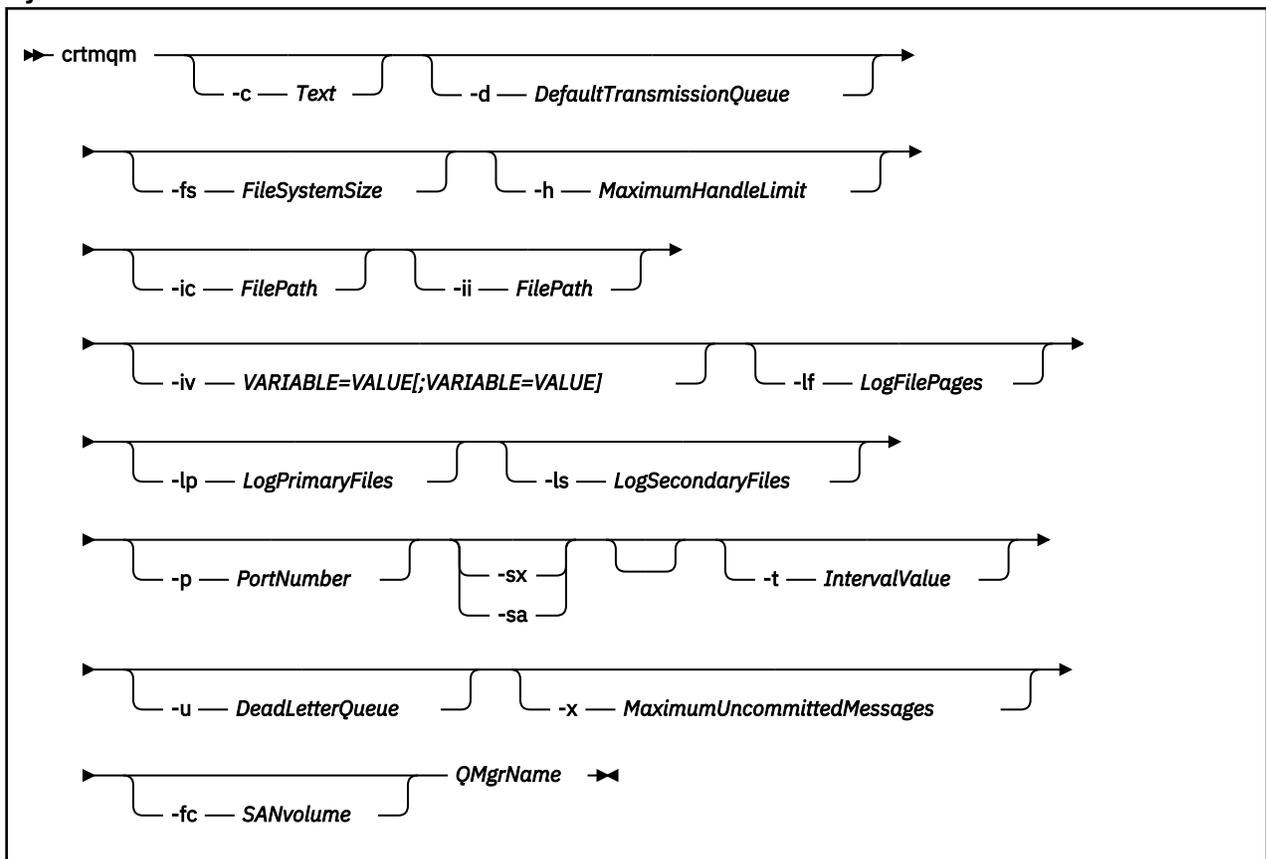
crtmqm (create queue manager)

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.

The *FileSystemSize* is allocated from the available disk space. A disaster recovery or high availability queue manager requires twice the disk space of a stand-alone queue manager.

-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.

V 9.1.4 -ic FilePath

Automatic configuration of MQSC attributes.

Specify a location containing MQSC commands to be automatically applied to the queue manager on every queue manager restart. You can specify a filename, or a directory where each file with the .mqsc suffix is automatically processed. The filepath must be under mqbackup: //. For information about this feature, see [Automatic configuration from an MQSC script at startup](#) in the IBM MQ documentation.

V 9.1.4 -ii FilePath

Automatic configuration of INI attributes.

Specify a location containing INI format information to be automatically applied as an override to the qm.ini file on every queue manager restart. You can specify a filename, or a directory where each file with the .ini suffix is automatically processed. The filepath must be under mqbackup: //. For information about this feature, see [Automatic configuration of qm.ini at startup](#) in the IBM MQ documentation.

V 9.1.4 -iv VARIABLE=VALUE[;VARIABLE=VALUE]

Configuration variable for use with automatic uniform clusters.

Specify a name and associated value for use as an insert when processing MQSC definitions. Currently only used for CONNAME fields in defining cluster receivers for automatic uniform clusters.

-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.

-sa

Automatic queue manager startup. The queue manager is configured to start automatically when the appliance restarts. This argument is mutually exclusive with -sx.

-sx

Specifies that the queue manager is a high availability (HA) queue manager.

The queue manager starts automatically as part of the HA group. This argument is mutually exclusive with -sa.

-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.

M 2000 M 2001 -fc *SANvolume*

Specifies that the queue manager uses SAN storage. The LUN that the queue manager is associated with is identified by a previously-created volume object specified by *SANvolume*. This option is mutually exclusive with the -sx option, because SAN storage is not available to high availability queue

managers. (SAN storage is available only on M2001 and M2000 appliances, it is not available on M2002 appliances.)

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 queue manager)

Delete a queue manager.

Purpose

You can use the **`dltmqm`** command to delete a queue manager.

Syntax

```
►► dltmqm — QMgrName ◄◄
```

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 (mqc1i) #`. To enter the IBM MQ administration mode, enter `mqc1i` 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)

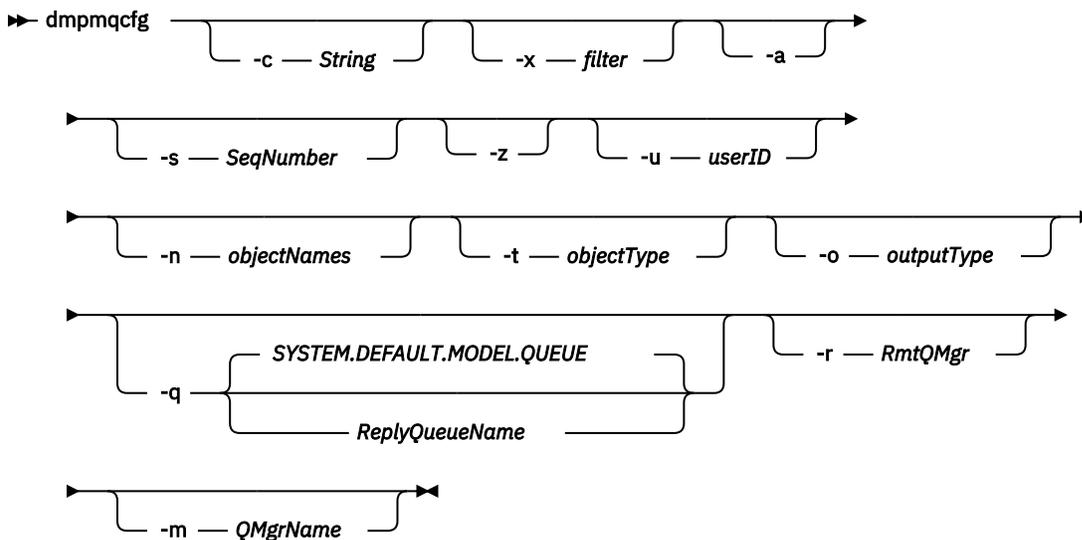
dmpmqcfg (dump queue manager configuration)

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('conname')"

Specifies that the specific client channel specified by *chlname* is used to connect to the queue manager at *conname*.

conname 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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 `dmpmqcfig` 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 `dmpmqcfig` 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 `dmpmqcfig`.
- 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
```

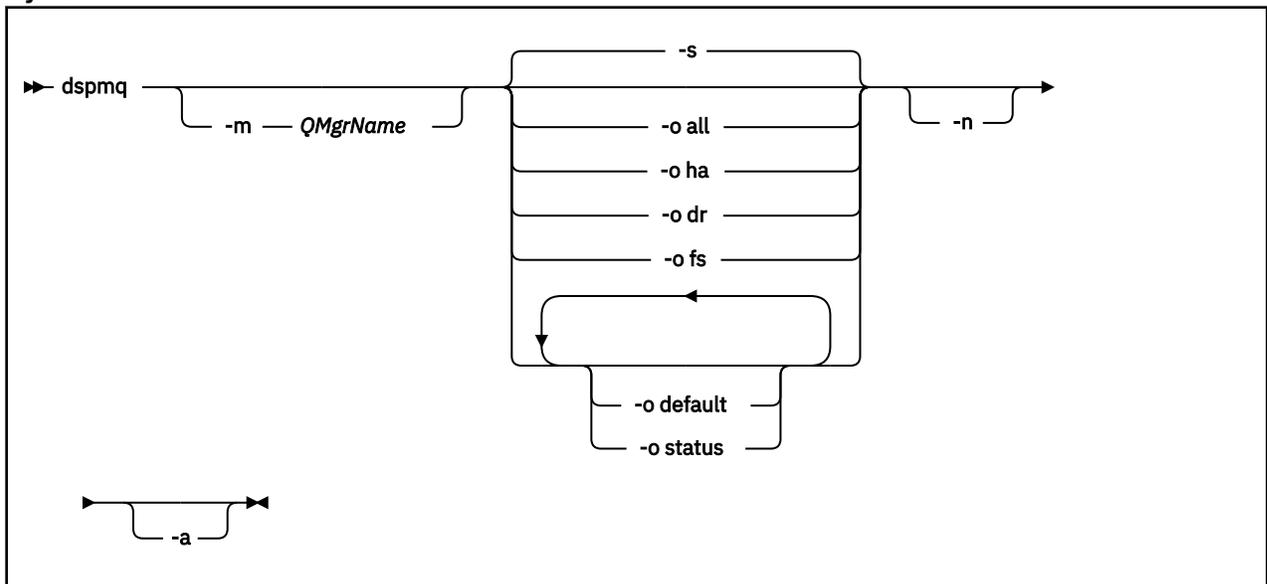
dspmq (display queue managers)

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 fs

Specifies that information about the queue manager file system is displayed. For a queue manager that uses SAN storage, it gives the volume label of the associated device. (SAN storage is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.)

-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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 (display route information)

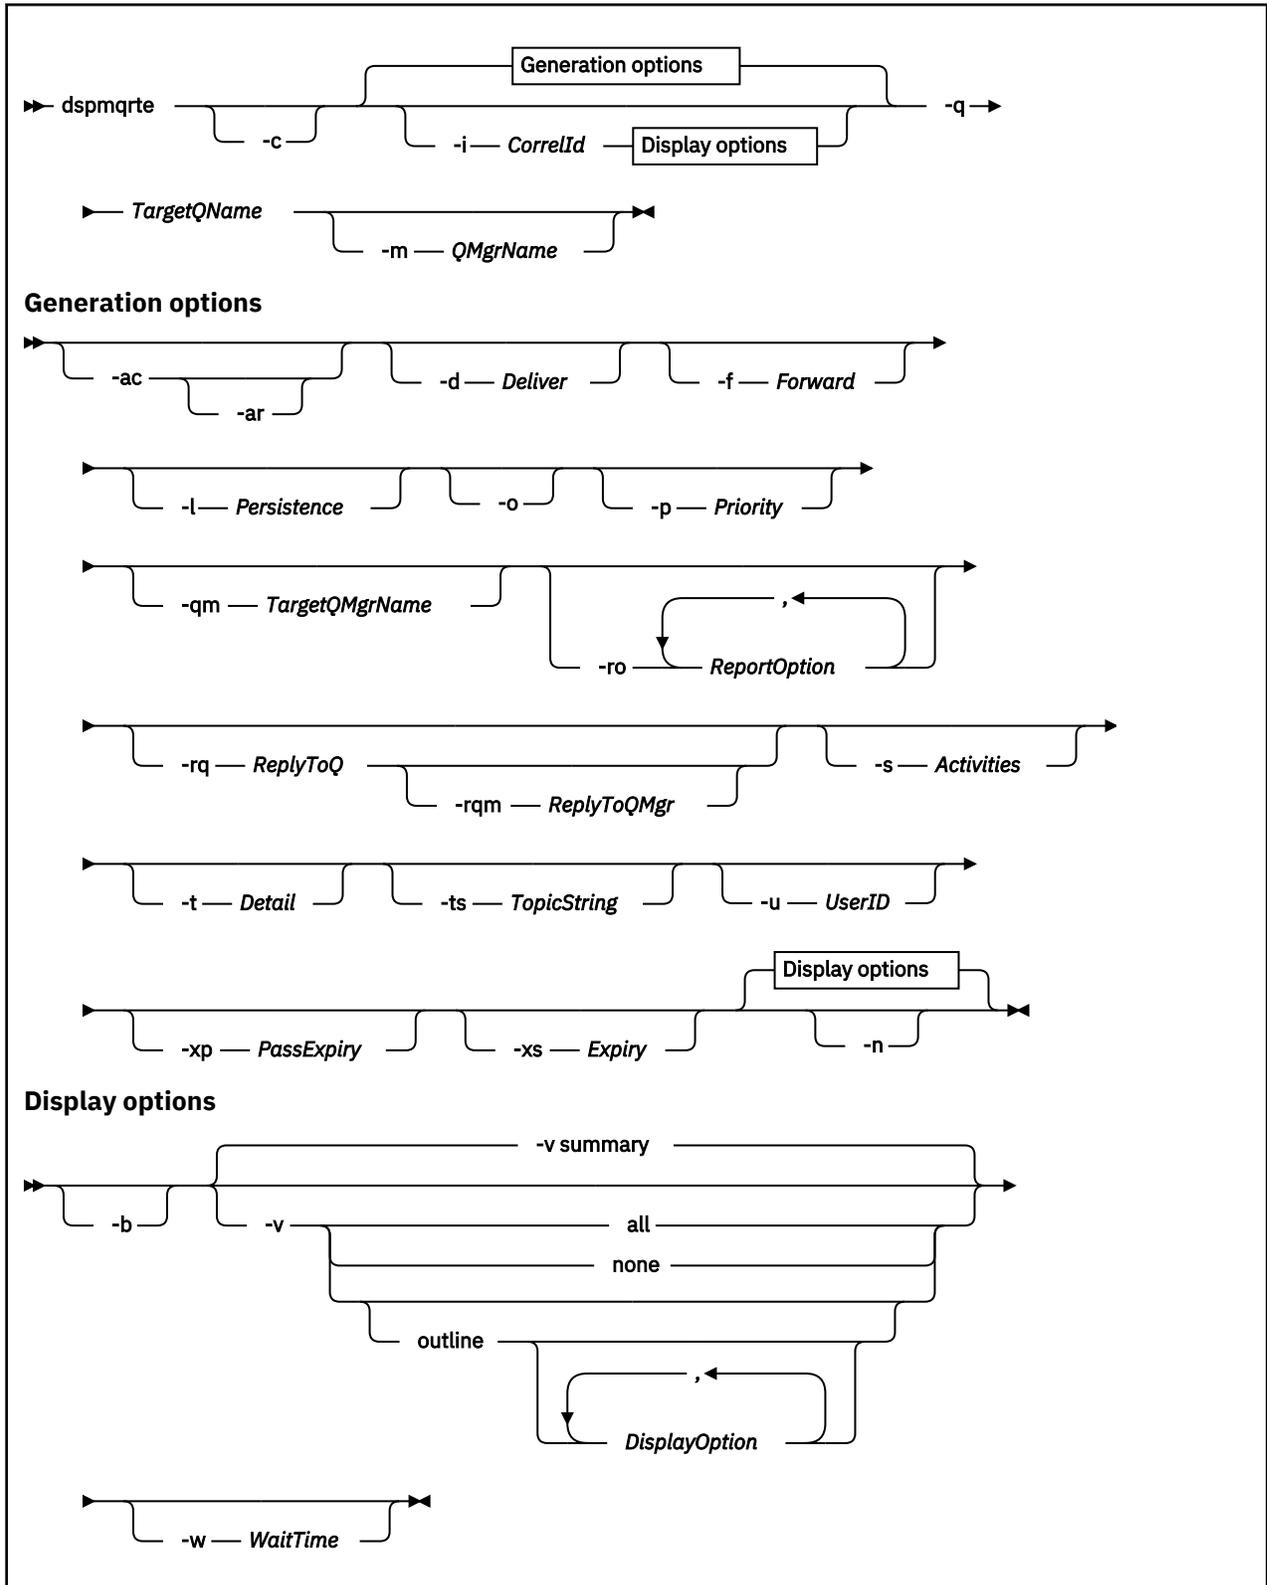
Determine the route that a message has taken through a queue manager network.

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



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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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.

V 9.1.5 *dspmqrtrc* (display formatted trace)

Format and display IBM MQ trace.

Purpose

Use the **dspmqrtrc** command to display IBM MQ formatted trace output.

Syntax



Required parameters

InputFileName

The name of the file containing the unformatted trace, for example:

```
AMQ12345.01.TRC
```

The output file is named AMQ *yyyyy.zz.FMT*, based on the PID of the trace file and is written to the `mqtrace://` directory.

Specify an asterisk character (*) to format all trace files.

Optional parameters

-h

Omit header information from the report.

-s

Extract trace header and put to stdout.

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Trace files are created in the `mqtrace://` directory.
- Formatted files are written to the `mqtrace://` directory.

Related commands

Table 51. Related command names and descriptions

Command	Description
“endmqtrc (end trace)” on page 599	End trace
“strmqtrc (Start trace)” on page 613	Start trace

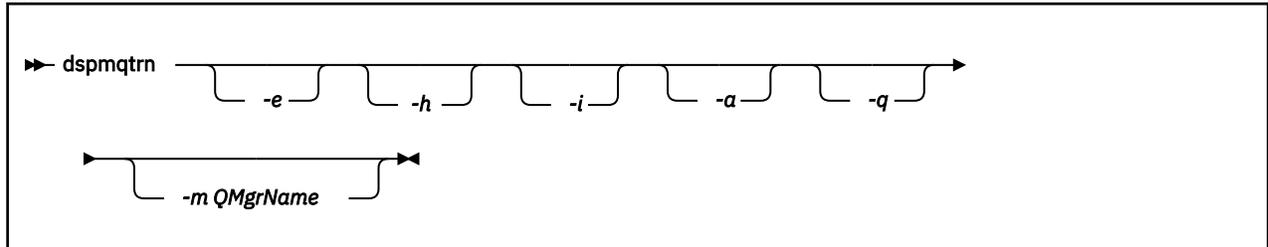
dspmqtrn (display incomplete transactions)

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(mqc1i)#`. 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](#)

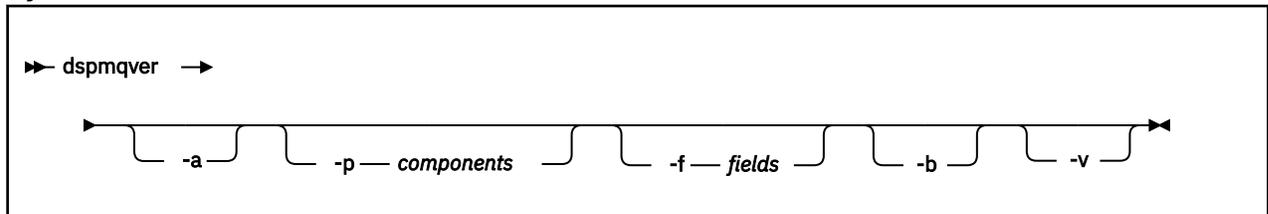
dspmqver (display version information)

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
```

dspmqweb (display mqweb server configuration)

Display information about the configuration of the mqweb server. The mqweb server is used to support the IBM MQ Console and administrative REST API.

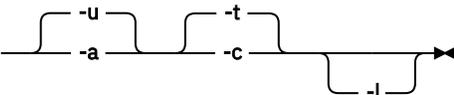
Purpose

Use the **dspmqweb properties** command to view details of the configuration of the mqweb server.

The following configuration properties are available on the appliance:

- mqRestRequestTimeout - REST request timeout
- traceSpec - Trace specification
- maxTraceFileSize - Maximum trace file size
- maxTraceFiles - Maximum number of trace files
- ltpaExpiration - LTPA token expiration
- mqRestCorsAllowedOrigins - REST CORS allowed origins
- mqRestCorsMaxAgeInSeconds - Maximum REST CORS age

Syntax

```
➔ dspmqweb — properties — 
```

Optional parameters

properties

Displays information about the configurable properties of the mqweb server. That is, which properties are configurable by the user and those that have been modified.

-u

Displays only the configurable properties that have been modified by the user.

-a

Displays all available configurable properties, including those which have been modified by the user.

-t

Formats the output as text name-value pairs.

-c

Formats the output as command text which can be used as input to the corresponding **setmqweb properties** command.

-l

Enable verbose logging. Diagnostic information is written to a mqweb server log-file.

Return codes

Return code	Description
0	Command successful
>0	Command not successful.

Related reference

“setmqweb (set mqweb server configuration)” on page 609
Add or remove a known mqweb server configuration property.

endmqm (end queue manager)

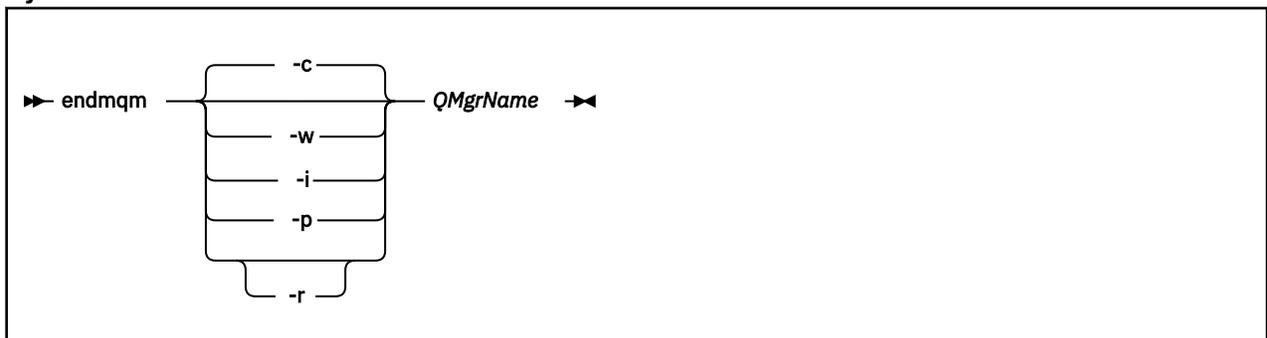
Stop a queue manager.

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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)
- [dlmqm](#) (Delete queue manager)

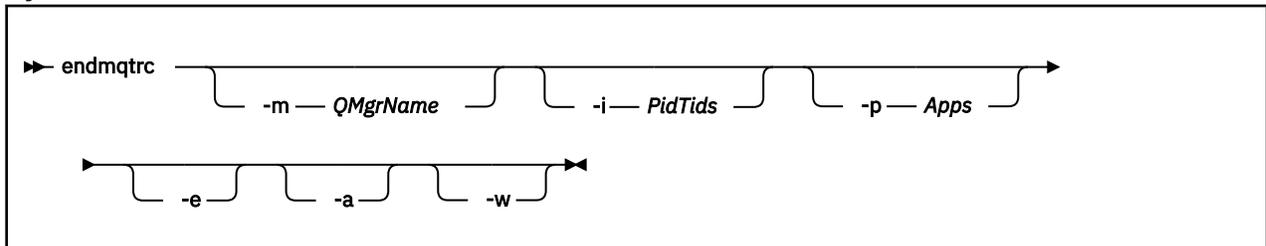
endmqtrc (end trace)

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(mqc.li)#`. To enter the IBM MQ administration mode, enter `mqc.li` 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 \(Start trace\)”](#) on page 613

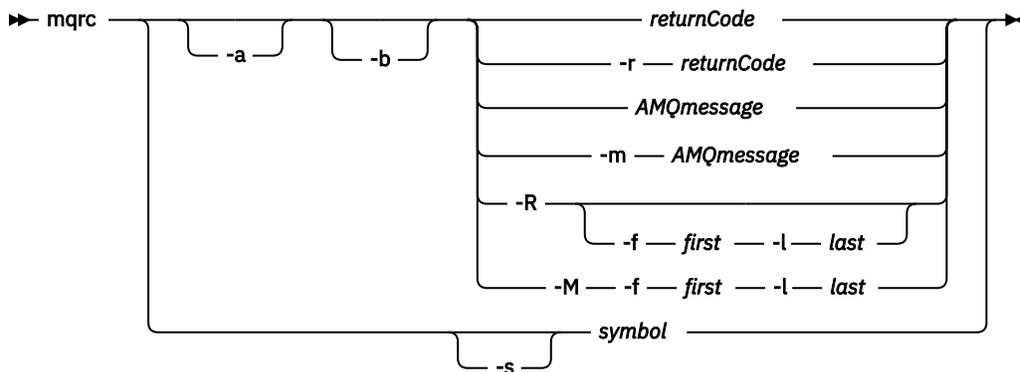
mqrc

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
```

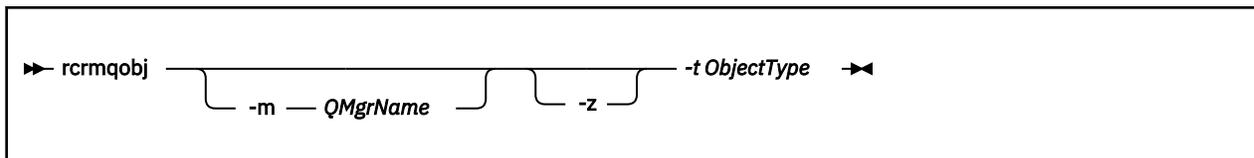
- This command displays return codes in the range 2505 - 2530:

```
mqrc -R -f 2505 -l 2530
```

rcrmqobj (re-create object)

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.

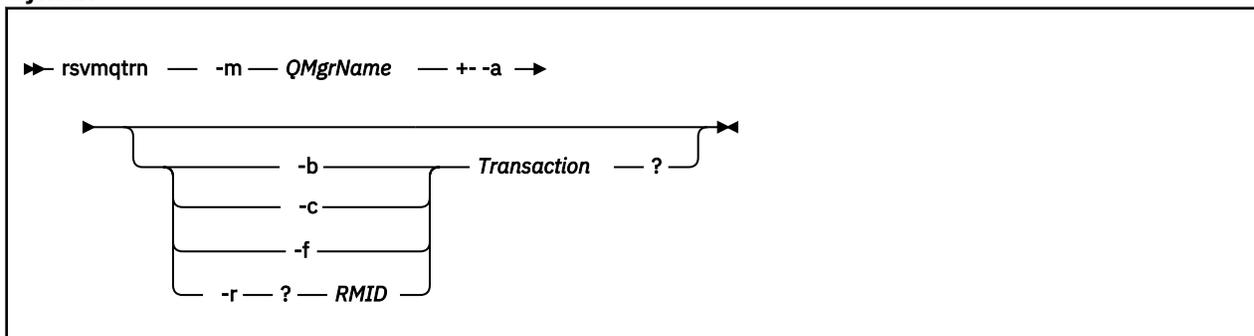
rsvmqtrn (resolve transactions)

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(mqc.li)#`. To enter the IBM MQ administration mode, enter `mqc.li` 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 \(display incomplete transactions\)” on page 594](#)

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.

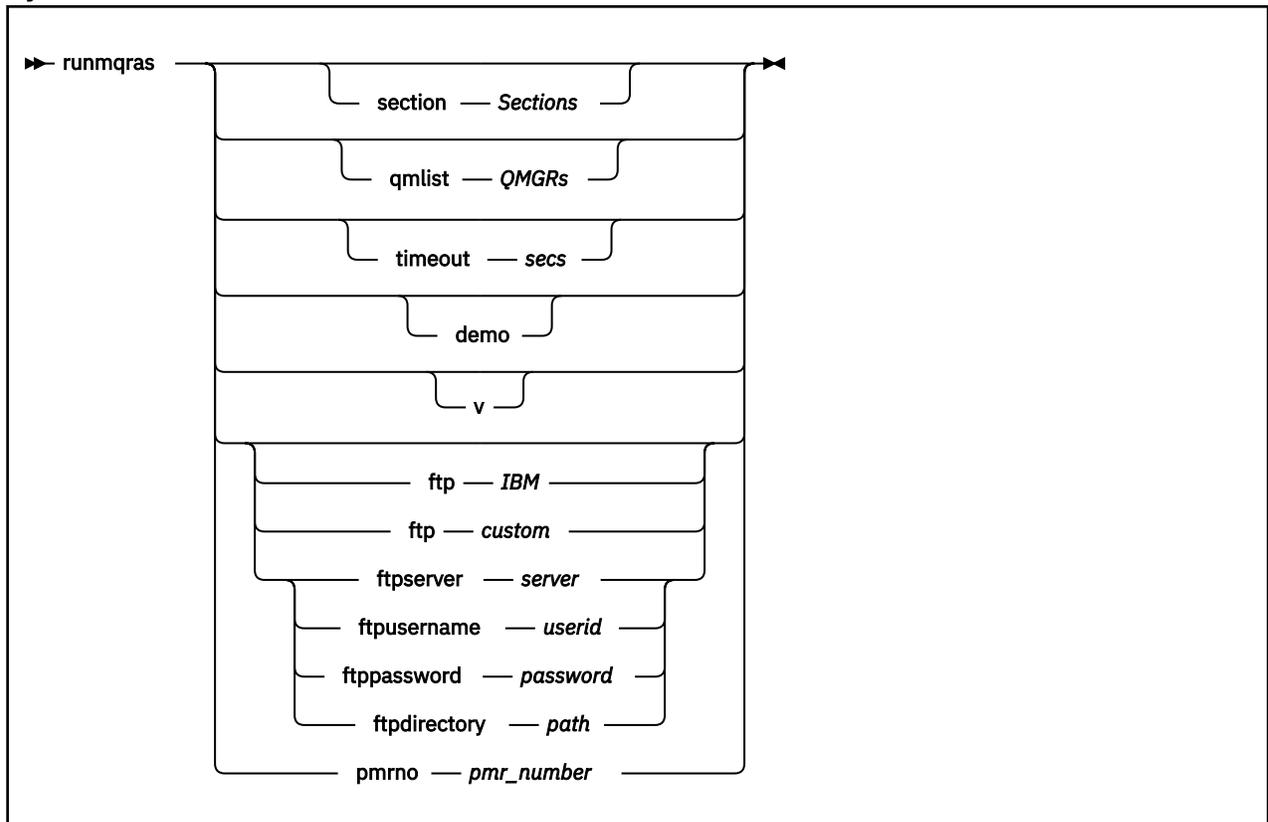
By default, the command gathers information such as the FDC files, error logs, product version, and status information. The command does not gather user information that is contained in messages on queues

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 788), or by using the IBM MQ Appliance web UI (see “Managing files by using the IBM MQ Appliance web UI” on page 402). 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.

You could also consider using the `qmlist` parameter with this option to limit the output to specified queue managers.

cluster

Gather information specific for clustering

defs

Gather the queue manager definitions and status

format

Generate and include formatted trace files.

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.
- Use [V 9.1.5](#) `runmqras -section format` to generate and include formatted trace files.
- 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
```

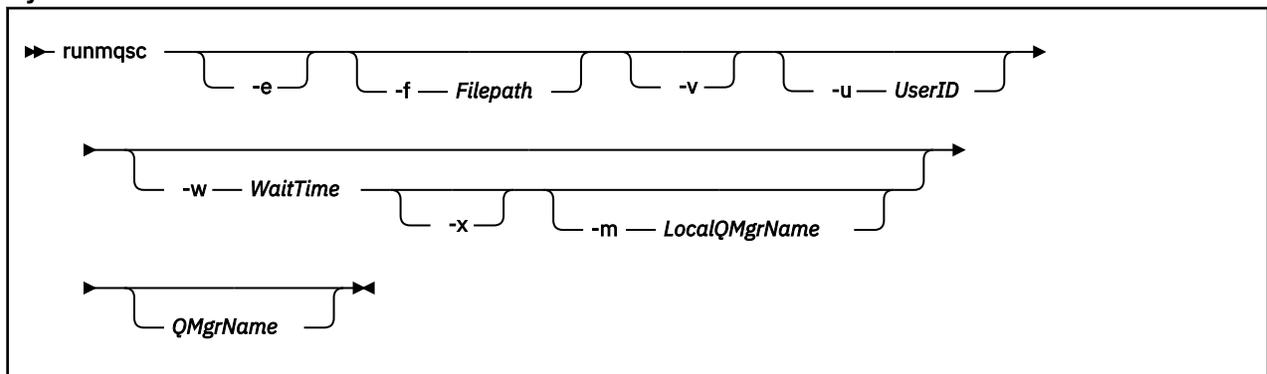
runmqsc (run MQSC commands)

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 9.1.4 -f filepath

Read the input to be processed from the specified file rather than standard input. The supplied filepath should be a uri in either `mqbackup://` or `mqtemporary://`.

-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.

-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 [Configuring queue managers 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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 on the appliance takes its input directly from the terminal, therefore you can enter MQSC commands interactively from the keyboard at the `mqc1i` prompt. In some cases, for example, automating a complex set of steps, you might want to run **runmqsc** remotely from a

distributed IBM MQ client. You can then place your commands in a file and redirect the input from that file (see `runmqsc`).

- When the commands are processed, the results and a summary are put into a report that is sent to your terminal. You might wish to capture this output to a file for reporting purposes, details on how to do this depend on your SSH client. Alternatively, you can run `runmqsc` from an IBM MQ client and redirect the output 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.

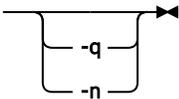
Purpose

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

```
runswchl -m QmgrName -c ChannelName [-q] [-n]
```



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 `mqc1i(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 `T0.QM1`:

```
runswchl -m QM1 -c T0.QM1 -q
```

- The following command switches the transmission queue for cluster-sender channel `T0.QM3` without moving the messages on it:

```
runswchl -m QM1 -c T0.QM3 -n
```

- The following command switches the transmission queue for cluster-sender channel `T0.QM3` and move the messages on it:

```
runswchl -m QM1 -c T0.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 `T0.*`:

```
runswchl -m QM1 -c T0.* -q
```

setmqweb (set mqweb server configuration)

Add or remove a known mqweb server configuration property.

Purpose

You can use the **setmqweb properties** command to configure the mqweb server. The mqweb server is used to support the IBM MQ Console and administrative REST API.

Syntax

```
➤ setmqweb — properties — -r — -k — name — -d — -v — value — -l
```

Parameters

-r

Reset to default values. This parameter removes all user-modified configuration properties .

-k *name*

The name of the configuration property to add, update or remove. The following list shows the valid values for *name* on the appliance:

- mqRestRequestTimeout
- traceSpec
- maxTraceFileSize
- maxTraceFiles
- ltpaExpiration
- mqRestCorsAllowedOrigins
- mqRestCorsMaxAgeInSeconds

-d

Deletes the specified configuration property.

-v *value*

The value of the configuration property to add to or update. Any existing configuration properties of the same *name* are overwritten. Duplicate configuration properties are removed. The value is case-sensitive and can be enclosed in double quotation marks to allow for multiple token or empty values. The *value* that is specified is not validated. If incorrect values are specified a subsequent attempt to start the mqweb server might fail. For possible values for each *name*, see [“Configuring the IBM MQ Console and REST API”](#) on page 306.

-l

Enable verbose logging. Diagnostic information is written to an mqweb server log-file.

Return codes

Return code	Description
0	Command successful
>0	Command not successful.

Related reference

[“dspmqweb \(display mqweb server configuration\)”](#) on page 596

Display information about the configuration of the mqweb server. The mqweb server is used to support the IBM MQ Console and administrative REST API.

strmqm (start queue manager)

Start a queue manager.

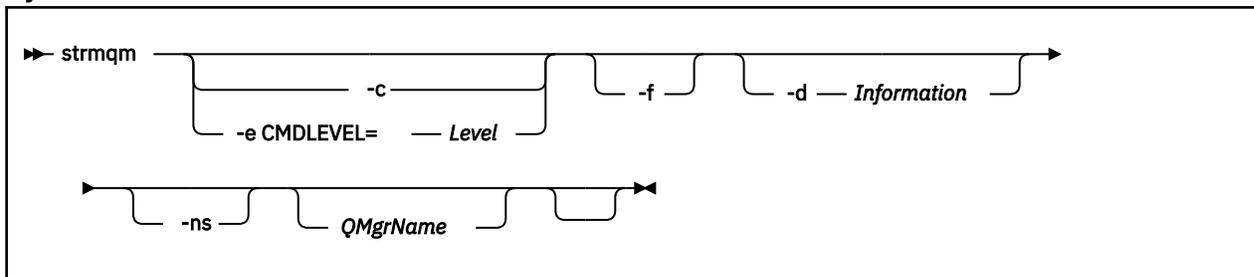
Purpose

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 881.

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 **mq strmqm** command again to start the queue manager.

If you run `mq strmqm -c` on a queue manager that is being used as an 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; you can use this option together with the `-c` option to reset the HA queue manager without starting it. (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 (mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 HA 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
```

- The following command resets the HA queue manager HAQM1 without starting it on the issuing appliance:

```
strmqm -cns HAQM1
```

Related commands

- [crtmqm](#) (Create queue manager)
- [endmqm](#) (Start queue manager)
- [dlmqm](#) (Delete queue manager)

strmqtrc (Start trace)

Start trace at a specified level of detail, or report the level of tracing in effect.

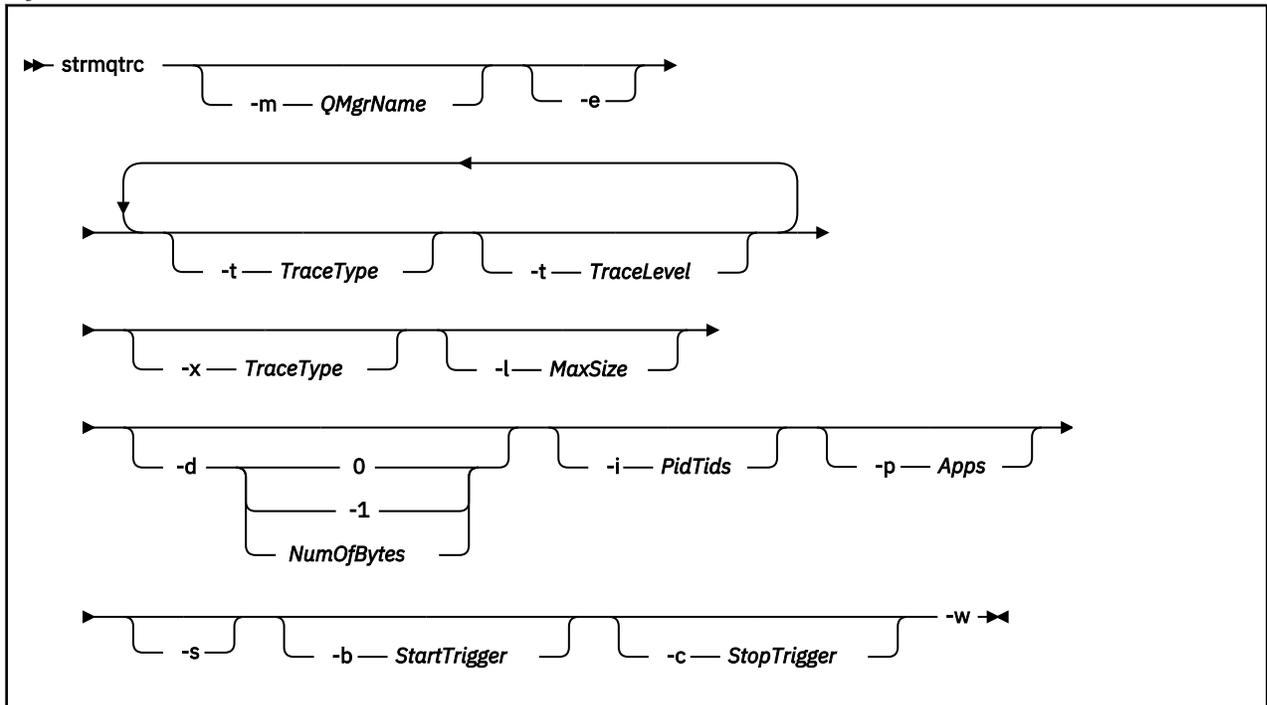
Purpose

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 788), or by using the IBM MQ Appliance web UI (see “Managing files by using the IBM MQ Appliance web UI” on page 402)

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(mqc.li)#`. To enter the IBM MQ administration mode, enter `mqc.li` 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 **enmqtrc** 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 `AMQppppp.qq.TRS` and a new `AMQppppp.qq.TRC` file is started. If a previous copy of an `AMQppppp.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 endmqtrc` 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 \(end trace\)” on page 599](#)

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`.

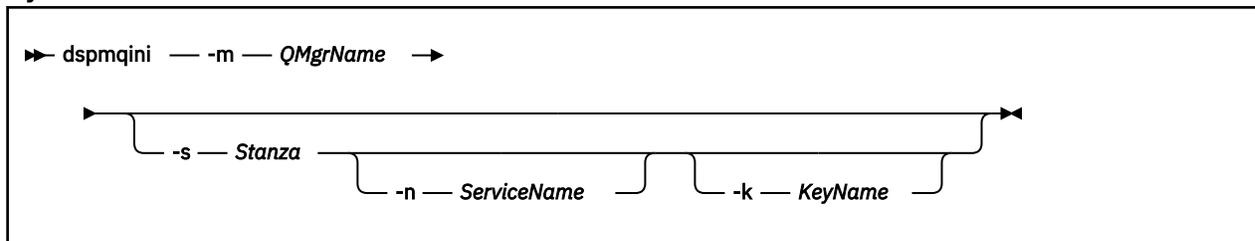
dspmqini (display mq ini)

Display attributes from the `qm.ini` or `mqat.ini` file of a specified queue manager. Display attributes from the `mqs.ini` file.

Purpose

You can use the `dspmqini` command to display the `qm.ini` or `mqat.ini` file for a queue manager. You can also view the `mqs.ini` file.

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
- DiagnosticMessages
- InstanceData
- TuningParameters
- SSL
- Security
- Subpool
- **V 9.1.4** AutoConfig
- **V 9.1.4** AutoCluster
- **V 9.1.4** Variables

The following value for *Stanza* is valid for the `mqat.ini` file:

- AllActivityTrace

The following values for *Stanza* are valid for the `mq.s.ini` file:

- DiagnosticSystemMessages
- DiagnosticMessagesTemplate

If you do not specify any further parameters, the command displays the values for the stanza specified.

-n *ServiceName*

For `DiagnosticSystemMessages`, `DiagnosticMessages`, and `DiagnosticMessagesTemplate` stanzas, specifies the name of the service being displayed.

-k *KeyName*

Specifies which key (attribute) of the stanza is displayed. (See [“setmqini \(set mq ini\)”](#) on page 620 for more information about valid keys.)

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(mqc1i)#`. 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
```

- The following command displays the DiagnosticMessages stanza of the service named Filtered that has been defined for queue manager QM1:

```
dspmqini -m QM1 -s DiagnosticMessages -n Filtered
```

Related commands

- [setmqini](#)

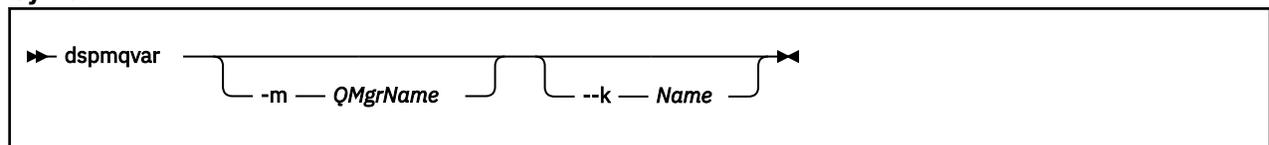
dspmqvar (display mq variables)

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`:

```
dspmqlvar -m QM1 -k AMQ_SERVICE_DEBUG_REPOS
```

- The following command displays the values of all the global environment variables:

```
dspmqlvar
```

Related commands

- [setmqvar](#)

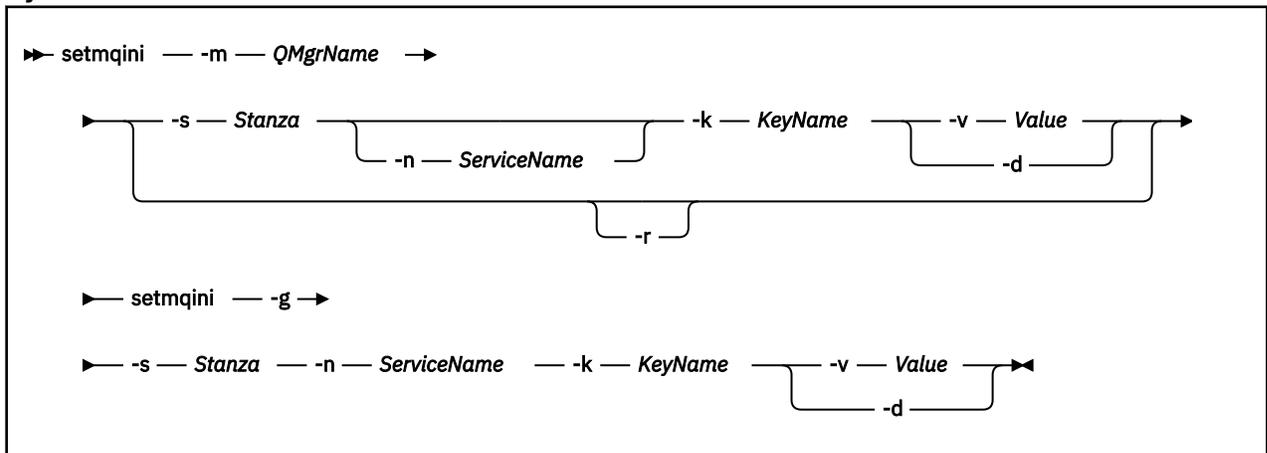
setmqini (set mq ini)

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. Set or delete a diagnostic message definition in the `mqs.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. You also use the command to define `DiagnosticSystemMessages` and `DiagnosticMessagesTemplate` services in the `mqs.ini` file.

Syntax



Parameters

-m *QMGrName*

Specifies that the `qm.ini` file that is associated with the specified queue manager is to be modified.

-g

Specifies that `mqs.ini` file is to be updated with details of a `DiagnosticSystemMessages` or a `DiagnosticMessagesTemplate` service.

-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
- DiagnosticMessages
- InstanceData
- TuningParameters
- SSL
- Security
- Subpool
- **V 9.1.4** AutoConfig
- **V 9.1.4** AutoCluster
- **V 9.1.4** Variables

The following value for *Stanza* modifies the `mqat.ini` file.

- AllActivityTrace

The following values for *stanza* modify the `mqs.ini` file

- DiagnosticSystemMessages
- DiagnosticMessagesTemplate

V 9.1.4 -r

Permanently resets the automatic configuration INI baseline (applicable when a queue manager is a member of a uniform cluster that uses automatic configuration).

When this option is used, the `base_qm.ini` file is removed from that queue manager. This is the file that is the baseline to which auto-configuration is reapplied on each queue manager start. Specifying the `-r` option causes the current queue manager configuration to become the baseline on future starts.

-n ServiceName

For `DiagnosticSystemMessages`, `DiagnosticMessages`, and `DiagnosticMessagesTemplate` stanzas, specifies the name of the service being defined or deleted.

-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.

The valid key names for the various stanzas can be retrieved from the IBM MQ documentation, by using the following links:

- Log - [Log defaults for IBM MQ](#)
- TCP - [TCP, LU62, and NETBIOS](#)
- Channels - [Channels stanza of the qm.ini file](#)
- DiagnosticMessages, DiagnosticSystemMessages, DiagnosticMessagesTemplate - [Diagnostic message service stanzas](#)
- InstanceData - [“Restarting queue managers by using the command line” on page 357](#)
- TuningParameters - [Tuning parameters](#)
- SSL - [SSL stanza of the queue manager configuration file](#)
- Security - [Security](#)
- Subpool - [Subpool](#)
- AllActivityTrace - [Configuring trace levels using mqat.ini](#)
- **V 9.1.4** AutoConfig - [Using automatic cluster configuration](#)

- [v 9.1.4](#) AutoCluster - [Using automatic cluster configuration](#)
- [v 9.1.4](#) Variable - [Using automatic cluster configuration](#)

-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. See the topics linked to from the `-k keyname` description for information about values.

-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. The exception to this rule is diagnostic message service stanzas. You can delete the entire stanza by specifying the `-n` parameter to identify the service whose stanza you want to delete.

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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](#).
- 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
```

- The following commands configure a `DiagnosticSystemMessages` service named `SystemErrors` that records messages with a severity of error, stop, or system, and writes them to files with the prefix 'SystemErrorsFile'.

```
setmqini -g -s DiagnosticSystemMessages -n SystemErrors -k Service -v File
setmqini -g -s DiagnosticSystemMessages -n SystemErrors -k Severities -v E+
setmqini -g -s DiagnosticSystemMessages -n SystemErrors -k FilePrefix -v SystemErrorsFile
```

- The following command deletes the stanza that defines the `DiagnosticSystemMessages` service named `SystemErrors`:

```
setmqini -g -s DiagnosticSystemMessages -n SystemErrors -d
```

Related commands

- [dspmqini](#)

setmqvar (set mq variable)

Add or remove an environment variable for the appliance or for a specified queue manager.

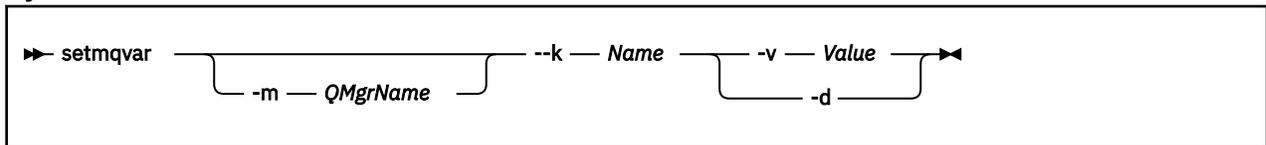
Purpose

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
```


Examples

- The following command modifies the user, `myid`, changes the password to be `newpassword`, and replaces the groups to which the user belongs to `adminUK` and `adminUS`:

```
usermodify -u myid -p newpassword -g MQadminUK, MQadminUS
```

userlist (list users)

Lists the messaging users on the IBM MQ Appliance.

Syntax

```
➤ userlist — -u — username — ➤
```

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 (create user group)

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

```
➤ groupcreate — -g — group — ➤
```

Parameters

-g *group*

Specifies the user group to be created.

Examples

- The following command creates a new user group, `MQgrp`:

```
groupcreate -g MQgrp
```

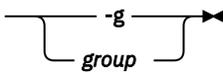
groupdelete (delete user group)

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

▶▶ groupdelete 

Parameters

-g group

Specifies the user group to be deleted.

Examples

- The following command deletes the user group, MQ0grp:

```
groupdelete -g MQ0grp
```

grouplist (list user groups)

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

▶▶ grouplist ▶▶

Examples

- The following command lists the user groups on the appliance:

```
grouplist
```

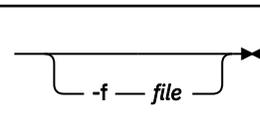
userbackup (back up users)

Backs up messaging users and groups 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 and groups that have been defined on the IBM MQ Appliance.

Syntax

▶▶ userbackup 

Parameters

-f filename

Specifies the file that the messaging users and groups 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 and groups to the file `backup_15115`:

```
userbackup -f backup_15115
```

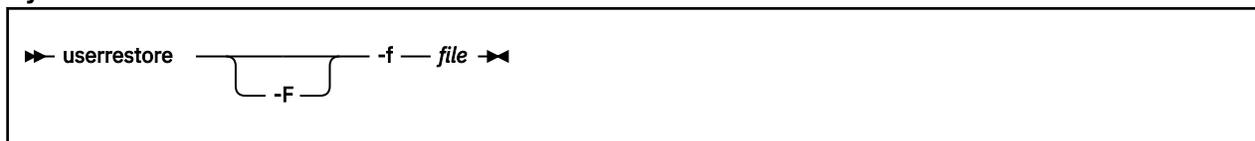
userrestore (restore user)

Restores messaging users and groups 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 and groups from back up file containing the details for messaging users and groups.

Syntax



Parameters

-F

If you specify the force option, the command replaces the configuration on the target appliance without requesting confirmation.

-f filename

Specifies the file containing the users and groups to restore. The file must be located in the `mqbackup:` location on the IBM MQ Appliance.

Usage notes

- All existing messaging user and group configurations in the target appliance are replaced by the user and group configurations from the backup file, with the exception of system users and groups, and the `mqm` user and group (which have special handling).
 - The system users include `root`, `mqm`, `mqsystem`, `hacluster`, and `sshd`.
 - The system groups include `root`, `mqm`, `haclient`, `sshd`, and `utmp`.
 - The membership of the `mqm` group on the target appliance will be replaced with the group membership of the `mqm` from the backup file. (The ID of the `mqm` group is never changed.)
- Any messaging users or groups that exist on the target appliance but are not present in the backup file are deleted from the target appliance.
- The **userrestore** command prompts for confirmation (Y/N) after processing the backup file, but before restoring the configuration. The number of users that will be deleted and added by the command is displayed.

Examples

- The following command restores users from the file backup_15115:

```
userrestore -f backup_15115

10 users and 5 groups will be added to the system, 5 users and
1 group will be deleted and mqm group membership will be modified.

Do you wish to proceed? Y/N
<Y>
userrestore completed successfully.
```

- The following command restores users from the file backup_15122, and does not ask for user confirmation:

```
userrestore -F -f backup_15122

8 users and 4 groups will be added to the system, 2 users and
1 group will be deleted and mqm group membership will be modified.

userrestore completed successfully.
```

Queue manager security management commands

You can use the certificate management commands to manage certificates for queue managers. If you have configured advanced message security for your queue managers, you can also implement MCA interception using security management commands.

The queue manager security commands can be run from the command line interface in MQ command mode. To enter MQ command mode, type `mqcli`.

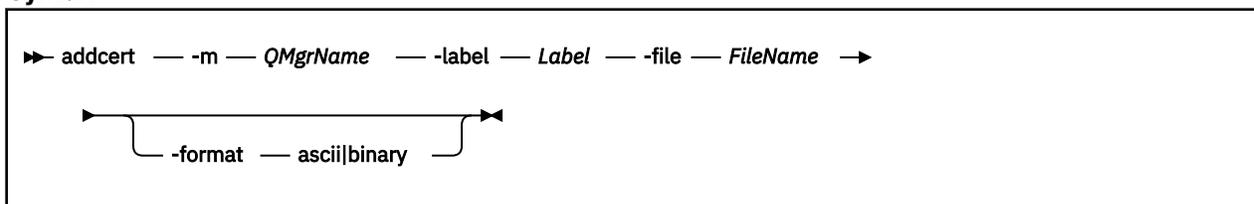
addcert

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 \(create certificate\)” on page 630](#)
- [“deletcert \(delete certificate\)” on page 634](#)
- [“detailcert \(detail certificate\)” on page 635](#)
- [“listcert \(list certificate\)” on page 639](#)
- [“receivecert \(receive certificate\)” on page 641](#)

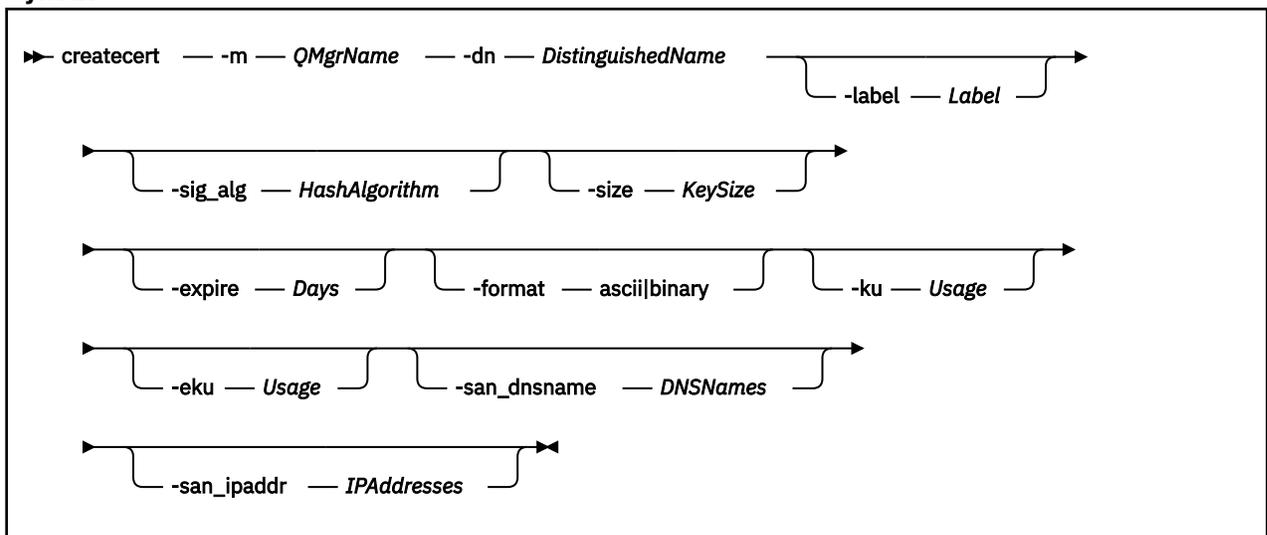
createcert (create certificate)

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 `ibmwebspheremqQMgrName`, 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 629](#)
- [“deletecert \(delete certificate\)” on page 634](#)
- [“detailcert \(detail certificate\)” on page 635](#)
- [“listcert \(list certificate\)” on page 639](#)
- [“receivecert \(receive certificate\)” on page 641](#)

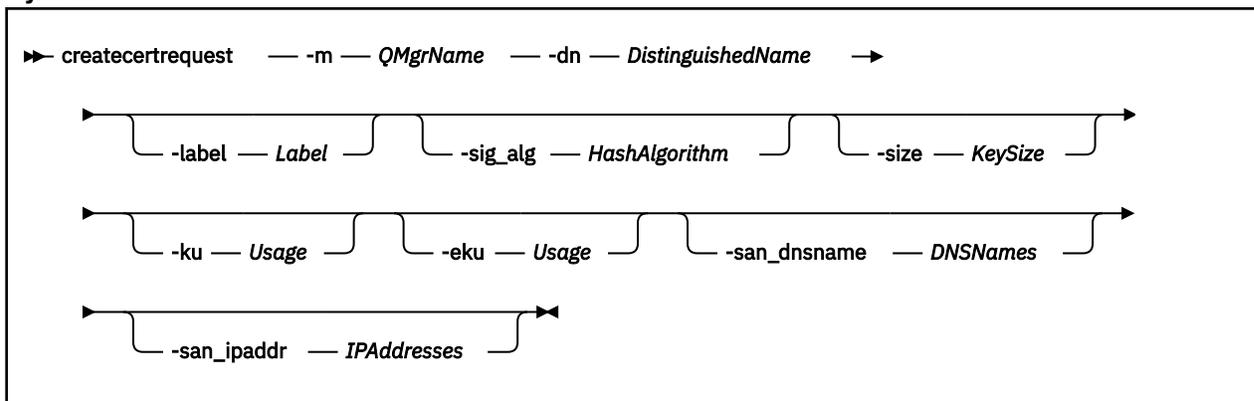
createcertrequest (create certificate request)

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 `ibmwebspheremqQMgrName`, 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 \(delete certificate request\)” on page 634](#)

- [“detailcertrequest \(detail certificate request\)” on page 636](#)
- [“listcertrequest \(list certificate request\)” on page 640](#)
- [“recreatecertrequest \(re-create certificate request\)” on page 642](#)

deletecert (delete certificate)

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

```
▶▶ deletecert — -m — QMgrName — -label — Label ▶▶
```

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 with a label of `cert1` for the queue manager `QM1`:

```
deletecert -m QM1 -label cert1
```

Related commands

- [“addcert” on page 629](#)
- [“createcert \(create certificate\)” on page 630](#)
- [“detailcert \(detail certificate\)” on page 635](#)
- [“listcert \(list certificate\)” on page 639](#)
- [“receivecert \(receive certificate\)” on page 641](#)

deletecertrequest (delete certificate request)

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.

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 `ibmwebspheremqmq1` for the queue manager `QM1` :

```
detailcert -m QM1 -label ibmwebspheremqmq1
```

Related commands

- [“addcert” on page 629](#)
- [“createcert \(create certificate\)” on page 630](#)
- [“deletecert \(delete certificate\)” on page 634](#)
- [“listcert \(list certificate\)” on page 639](#)
- [“receivecert \(receive certificate\)” on page 641](#)

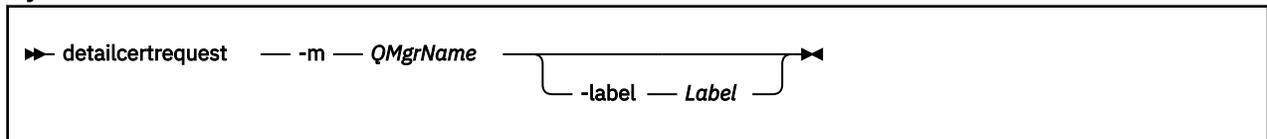
detailcertrequest (detail certificate request)

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 \(create certificate request\)” on page 632](#)
- [“deletecertrequest \(delete certificate request\)” on page 634](#)
- [“listcertrequest \(list certificate request\)” on page 640](#)
- [“recreatecertrequest \(re-create certificate request\)” on page 642](#)

dspamschl (set AMS channel)

Display information about MCA interception on queue managers.

Purpose

You can use the **dspamschl** command to display information about which queue managers have MCA interceptions set on one or more channels.

Syntax

```
▶▶ dspamschl -m QMgrName -n channel_name
```

Parameters

-m *QMgrName*

Specifies the name of the queue manager which you want to display MCA interception information for.

-n *channel_name*

Optionally specify a channel whose MCA interception status you are interested in.

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 queries MCA interception status for the server-connection channel SC1 on queue manager QM1:

```
dspamschl -m QM1 -n SC1
```

Related reference

[“setamschl \(set AMS channel\)” on page 643](#)

Set up MCA interception on a specific server-connection channel on a queue manager previously configured with AMS.

keybackup

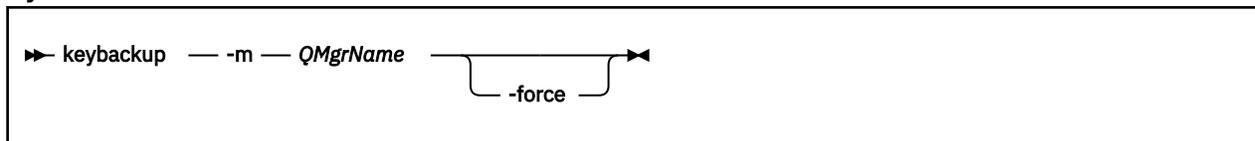
Back up the queue manager key repository to a file.

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 repository.

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” on page 638](#)
- [“copy” on page 788](#)

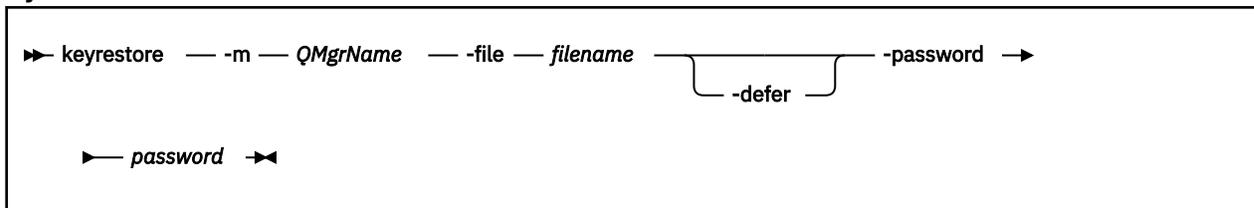
keyrestore

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 637](#)
- [“copy” on page 788](#)

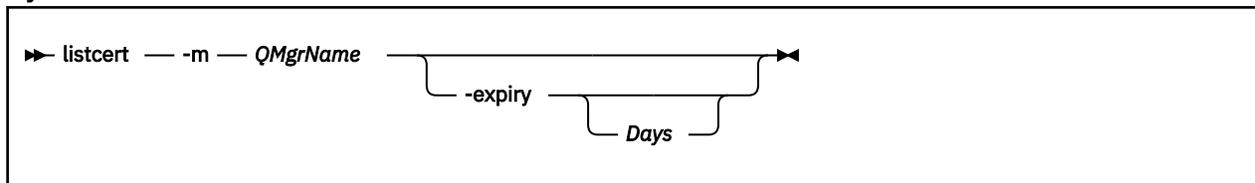
listcert (list certificate)

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 629](#)
- [“createcert \(create certificate\)” on page 630](#)
- [“deletecert \(delete certificate\)” on page 634](#)
- [“detailcert \(detail certificate\)” on page 635](#)
- [“receivecert \(receive certificate\)” on page 641](#)

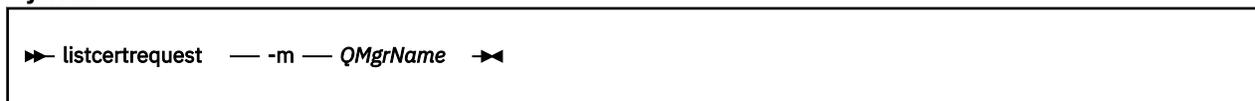
listcertrequest (list certificate request)

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 outstanding 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 \(create certificate request\)” on page 632](#)
- [“deletecertrequest \(delete certificate request\)” on page 634](#)
- [“detailcertrequest \(detail certificate request\)” on page 636](#)
- [“recreatecertrequest \(re-create certificate request\)” on page 642](#)

receivecert (receive certificate)

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

```
► receivecert — -m — QMgrName — -file — FileName — -format — ascii|binary ◄
```

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 629
- “[createcert \(create certificate\)](#)” on page 630
- “[deletecert \(delete certificate\)](#)” on page 634
- “[detailcert \(detail certificate\)](#)” on page 635
- “[listcert \(list certificate\)](#)” on page 639

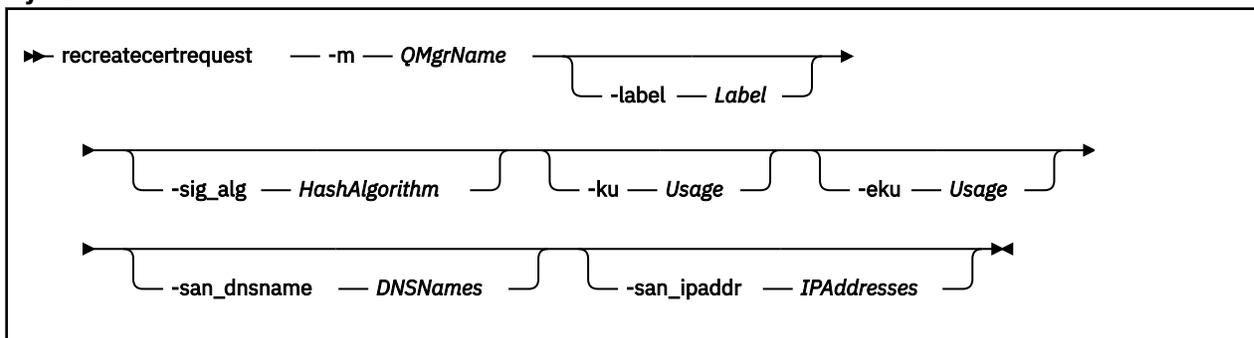
recreatecertrequest (re-create certificate request)

Re-create a certificate request for a specific queue manager.

Purpose

You can use the **recreatecertrequest** command to re-create 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 `ibmwebspheremqQMgrName`, 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 re-creates a certificate request for the queue manager QM1 :

```
recreatecertrequest -m QM1
```

Related commands

- [“createcertrequest \(create certificate request\)” on page 632](#)
- [“deletecertrequest \(delete certificate request\)” on page 634](#)
- [“detailcertrequest \(detail certificate request\)” on page 636](#)
- [“listcertrequest \(list certificate request\)” on page 640](#)

setamschl (set AMS channel)

Set up MCA interception on a specific server-connection channel on a queue manager previously configured with AMS.

Purpose

You can use the **setamschl** command to set up MCA interception on a particular server-connection channel on a specified queue manager. You can also use **setamschl** to delete existing MCA interceptions.

Syntax

```
►► setamschl — -m — QMgrName — -n — channel_name — -c — certificate_label —►
```

```
└───┬───┘  
    -d
```

Parameters

-m QMgrName

Specifies the name of the queue manager for which the MCA interception is required.

The queue manager must exist.

-n channel_name

Specifies the name of the server-connection channel for which the MCA interception is required.

-c *certificate_label*

Specifies the certificate used for the queue manager. The certificate is identified by its label.

-d

Specify this option to delete the specified MCA interception.

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.
- MCA interception requires that the required certificates and keys are held in the queue manager key repository. The key repository is created automatically with the queue manager and managed with the certificate commands.

Examples

- The following command creates an MCA interception for the server-connection channel SC1 on queue manager QM1:

```
setamschl -m QM1 -n SC1 -c cert1
```

Related reference

[“dspamschl \(set AMS channel\)” on page 637](#)

Display information about MCA interception on queue managers.

[“addcert” on page 629](#)

Add the public part of a certificate to the keystore of a specific queue manager.

[“createcert \(create certificate\)” on page 630](#)

Create a self-signed certificate for a queue manager.

[“deletecert \(delete certificate\)” on page 634](#)

Delete a certificate from the keystore of a specific queue manager.

[“listcert \(list certificate\)” on page 639](#)

List the certificates that are held in the keystore of a specific queue manager.

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 31](#) for specific information about using commands on the appliance.

M 2000 M 2001 [addmqm \(add existing queue manager\)](#)

Add an existing queue manager that uses SAN storage.

Purpose

You can use the **addmqm** command to re-create queue manager and reattach it to its SAN storage. You are most likely to need this command where the queue manager was previously running on an appliance that has failed.

Syntax

```
►► addmqm — -fc — SANvolume — -m — QMname ◄◄
```

Parameters

-fc SANvolume

Specifies the volume object that identifies the LUN that this queue manager uses.

This parameter is required.

-m QMName

Specifies the name of the queue manager that you want to re-create.

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

Examples

- The following command re-creates the queue manager QM1 using the volume SAN_QM1.

```
addmqm -fc SAN_QM1 -m QM1
```

Related commands

- [crtmqm](#) (Create queue manager)

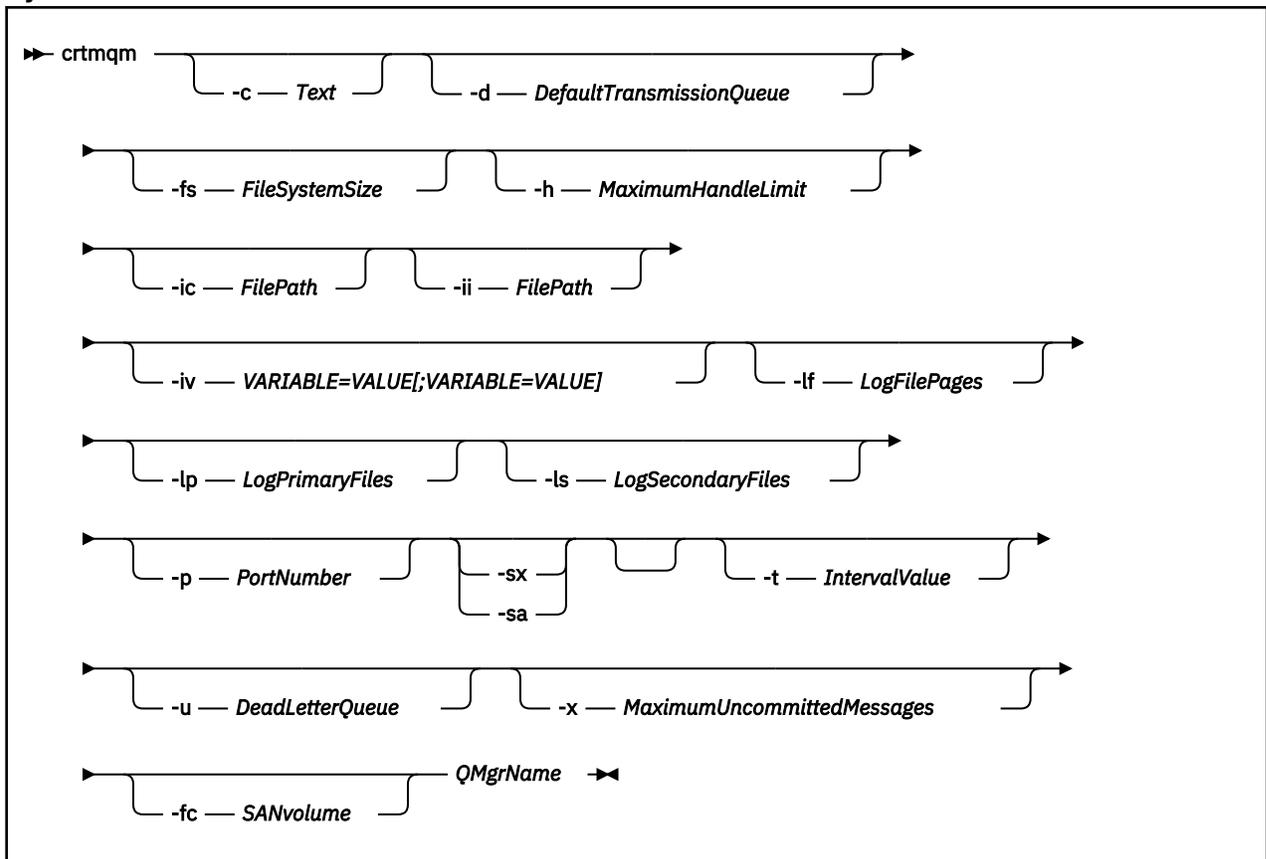
[crtmqm](#) (create queue manager)

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.

The *FileSystemSize* is allocated from the available disk space. A disaster recovery or high availability queue manager requires twice the disk space of a stand-alone queue manager.

-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.

V 9.1.4 -ic *FilePath*

Automatic configuration of MQSC attributes.

Specify a location containing MQSC commands to be automatically applied to the queue manager on every queue manager restart. You can specify a filename, or a directory where each file with the `.mqsc` suffix is automatically processed. The filepath must be under `mqbackup://`. For information about this feature, see [Automatic configuration from an MQSC script at startup](#) in the IBM MQ documentation.

V 9.1.4 -ii *FilePath*

Automatic configuration of INI attributes.

Specify a location containing INI format information to be automatically applied as an override to the `qm.ini` file on every queue manager restart. You can specify a filename, or a directory where each file with the `.ini` suffix is automatically processed. The filepath must be under `mqbackup://`. For information about this feature, see [Automatic configuration of qm.ini at startup](#) in the IBM MQ documentation.

V 9.1.4 -iv *VARIABLE=VALUE[;VARIABLE=VALUE]*

Configuration variable for use with automatic uniform clusters.

Specify a name and associated value for use as an insert when processing MQSC definitions. Currently only used for CONNAME fields in defining cluster receivers for automatic uniform clusters.

-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.

-sa

Automatic queue manager startup. The queue manager is configured to start automatically when the appliance restarts. This argument is mutually exclusive with -sx.

-sx

Specifies that the queue manager is a high availability (HA) queue manager.

The queue manager starts automatically as part of the HA group. This argument is mutually exclusive with -sa.

-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.

M 2000 M 2001 -fc SANvolume

Specifies that the queue manager uses SAN storage. The LUN that the queue manager is associated with is identified by a previously-created volume object specified by *SANvolume*. This option is mutually exclusive with the -sx option, because SAN storage is not available to high availability queue managers. (SAN storage is available only on M2001 and M2000 appliances, it is not available on M2002 appliances.)

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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:

```
ctrmqm -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.

```
ctrmqm -t 5000 -u SYSTEM.DEAD.LETTER.QUEUE QM2
```

Related commands

- [strmqm](#) (Start queue manager)
- [endmqm](#) (End queue manager)
- [dlmqm](#) (Delete queue manager)

dlmqm (delete queue manager)

Delete a queue manager.

Purpose

You can use the **dlmqm** command to delete a queue manager.

Syntax

```
▶▶ dlmqm — QMgrName ▶▶
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 [dlmqm](#) 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)

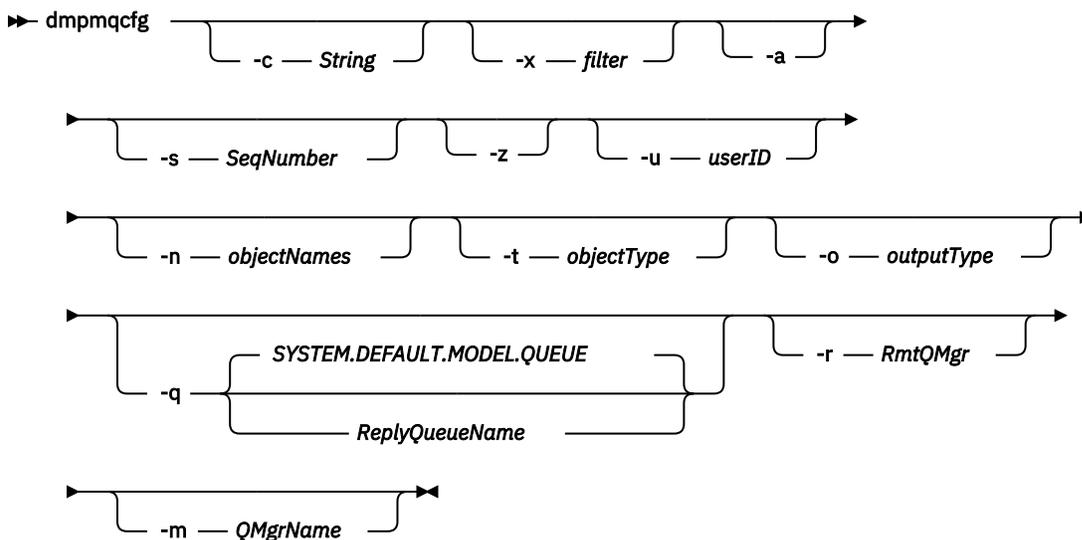
dmpmqcfg (dump queue manager configuration)

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('conname')"

Specifies that the specific client channel specified by *chlname* is used to connect to the queue manager at *conname*.

conname 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 `mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 **dmpmqcfig** 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 **dmpmqcfig** 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 **dmpmqcfig**.
- 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 [dmpmqcfig](#) in the IBM MQ documentation.

Examples

- The following command dumps the queue manager configuration for a queue manager `QM1`:

```
dmpmqcfig -m QM1
```

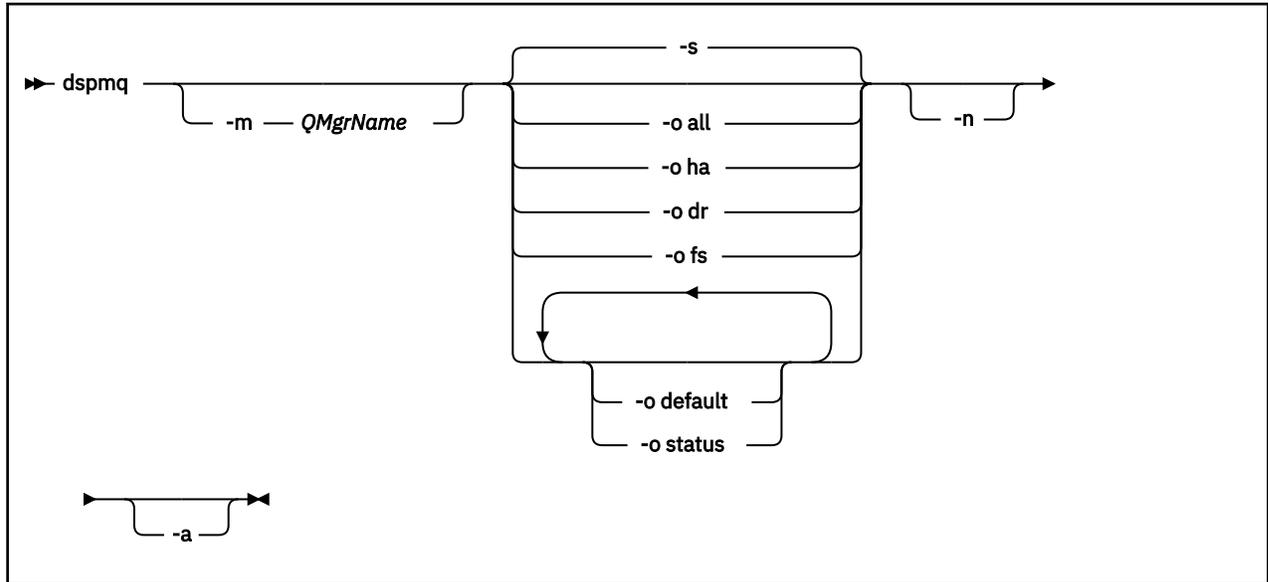
dspmq (display queue managers)

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 fs

Specifies that information about the queue manager file system is displayed. For a queue manager that uses SAN storage, it gives the volume label of the associated device. (SAN storage is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.)

-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 (end queue manager)

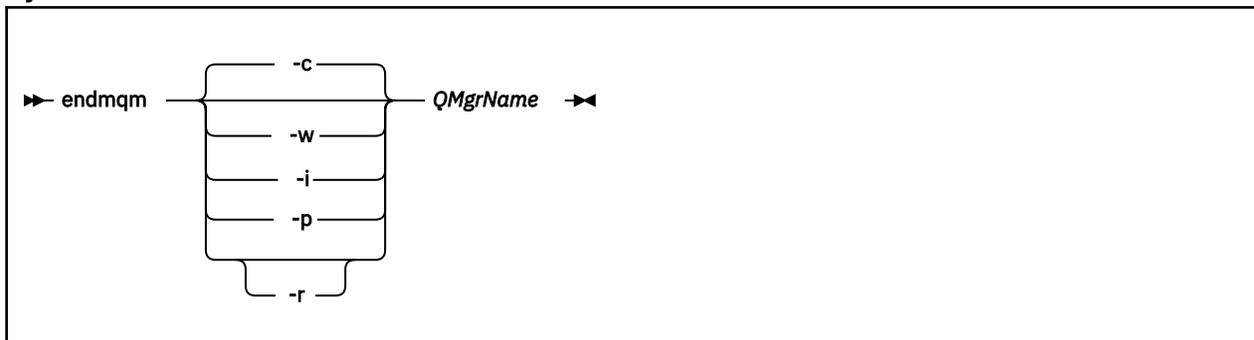
Stop a queue manager.

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 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 `mqccli)#`. To enter the IBM MQ administration mode, enter `mqccli` 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)
- [dlmqm](#) (Delete queue manager)

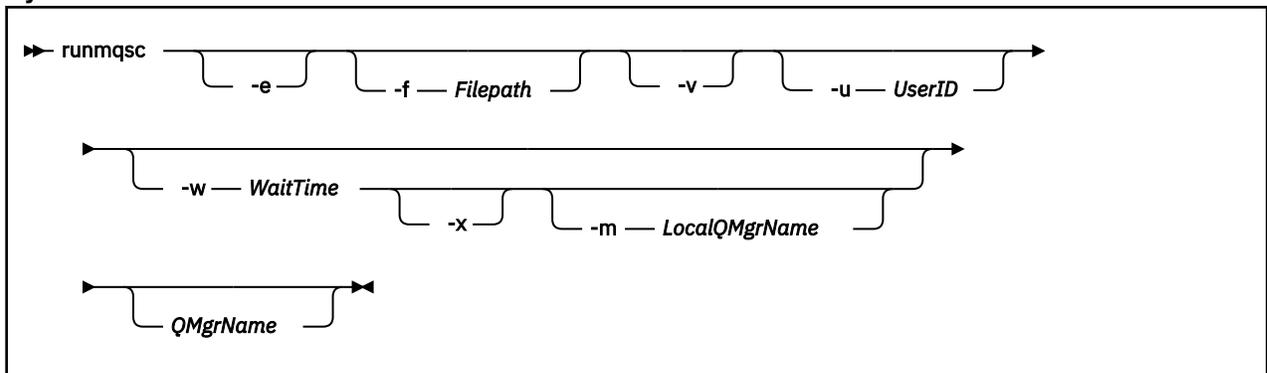
runmqsc (run MQSC commands)

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 9.1.4 -f filepath

Read the input to be processed from the specified file rather than standard input. The supplied filepath should be a uri in either `mqbackup://` or `mqtemporary://`.

-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.

-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 [Configuring queue managers 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 on the appliance takes its input directly from the terminal, therefore you can enter MQSC commands interactively from the keyboard at the `mqcli` prompt. In some cases, for example, automating a complex set of steps, you might want to run **runmqsc** remotely from a distributed IBM MQ client. You can then place your commands in a file and redirect the input from that file (see [runmqsc](#)).
- When the commands are processed, the results and a summary are put into a report that is sent to your terminal. You might wish to capture this output to a file for reporting purposes, details on how to do this

depend on your SSH client. Alternatively, you can run **runmqsc** from an IBM MQ client and redirect the output 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.

M2000 M2001 *rmvmqinf (remove queue manager that uses SAN)*

Remove a queue manager that uses SAN storage.

Purpose

You can use the **rmvmqinf** command to remove a queue manager that uses SAN storage

Note: SAN storage is available only on M2001 and M2000 appliances, it is not available for M2002 appliances.

Syntax

```
► rmvmqinf — QMname ◄
```

Parameters

QMName

Specifies the name of the queue manager that you want to remove.

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.
- Use this command to remove a queue manager that uses SAN storage. Do not use it to remove a queue manager that uses local SAN storage.

Examples

- The following command removes the queue manager QM1.

```
rmvmqinf QM1
```

Related commands

- [“addmqm \(add existing queue manager\)” on page 644](#)

setmqsize (increase queue manager file system size)

Increase the size of the file system allocated to a queue manager.

Purpose

When a queue manager is created, it is allocated file system space. This is 64 GB by default, but you can specify a different value if required when you create the queue manager. If you subsequently require a larger file system for that queue manager, you can expand it by using the **setmqsize** command.

Syntax

```
▶▶ setmqsize — -m — QMname — -s — size ▶▶
```

Parameters

-m QMName

Specifies the name of the queue manager that you want to expand the file system for.

This parameter is required.

-s size

Specifies the new size of the file system. Specify a positive integer with an optional M or G suffix to indicate megabytes or gigabytes. The value is taken as gigabytes if you do not specify otherwise.

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.
- This command cannot be run on a queue manager that uses SAN storage. (SAN storage is only available on M2000 or M2001 appliances, it is not available on M2002 appliances.)
- This command cannot be run on a queue manager that belongs to a high availability (HA) or disaster recovery (DR) configuration. If you require to expand the file system of an HA or DR queue manager, you must remove it from the HA or DR configuration, expand the file system size, then re-add it to the HA or DR configuration.
- The queue manager must be stopped before the command is run.
- The new size of the file system must be greater than or equal to the current file system size.
- Resizing file space for a queue manager does involve some I/O, and might degrade the performance of other queue managers while the resize is in progress.

Examples

- The following command expands the file storage for the queue manager QM1 to 128 GB.

```
setmqsize -m QM1 -s 128G
```

Related commands

- [crtmqm \(Create queue manager\)](#)

strmqm (start queue manager)

Start a queue manager.

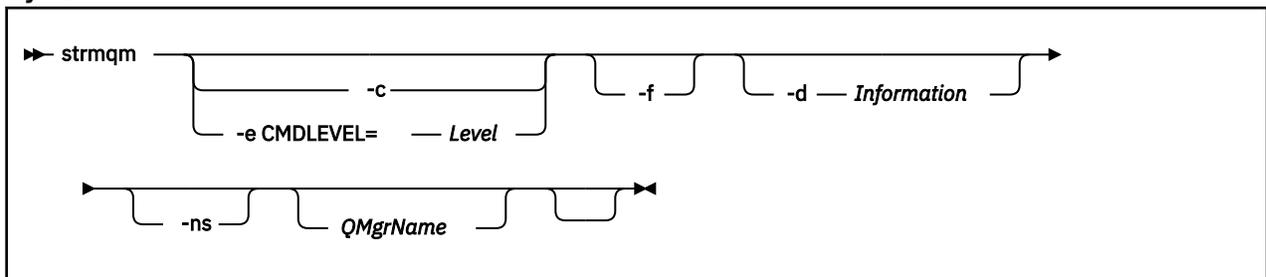
Purpose

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 881](#).

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 an 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; you can use this option together with the `-c` option to reset the HA queue manager without starting it. (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(mqc.li)#`. To enter the IBM MQ administration mode, enter `mqc.li` 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 HA 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
```

- The following command resets the HA queue manager HAQM1 without starting it on the issuing appliance:

```
strmqm -cns HAQM1
```

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.

This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `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 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.

Remote appliance(s) unavailable

This status means that the replication connection to the other appliance has been lost.

Partitioned

Queue manager data on the appliances is out of step, and cannot be automatically resolved.

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.

Inconsistent

The status is displayed on a secondary appliance during the initial synchronization of a queue manager if connection has been lost and synchronization was interrupted. The secondary appliance cannot provide high availability functionality until the initial synchronization has completed.

Resource failed

The physical interface associated with the floating IP used by the queue manager has been reconfigured. The queue manager is stopped and cannot be restarted until the original configuration is restored. This can also apply if an aggregated interface associated with the floating IP is changed.

This situation can also arise if the queue manager has been configured for disaster recovery and the eth20 interface has been reconfigured.

To recover, you must restore the original interface configuration, or remove the disaster recovery configuration. You can also delete the queue manager.

V 9.1.0 Migration Pending

When you have upgraded the appliance from Version 8.0 or Version 9.0 to Version 9.1, high availability queue managers display this status until you run the `sethagrp -u` command to upgrade them, and the other appliance is upgraded to Version 9.1 too. .

- 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 last synchronized. This information is shown when the status is `Partitioned` or `Inconsistent`.
- The time and date that the data was last in sync. This is displayed whenever data is out of sync.

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 an in-progress synchronization has been interrupted. 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.

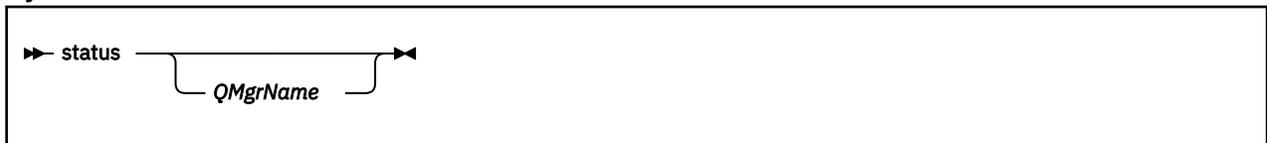
Remote appliance(s) not configured

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 last synchronized. This information is shown only when the status is `Partitioned`, `Inconsistent`, or `Remote appliance(s) unavailable`.
- The time and date that the data was last in sync. This is displayed whenever data is out of sync.
- 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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 for 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
```

Back up and restore commands

You can use the IBM MQback up and restore commands to back up queue managers, together with their log files and data.

The commands can be run from the command line interface in MQ command mode. To enter MQ command mode, type `mqcli`.

createbackupfs: create appliance storage location for back up

Allocate space for back up archive files on the appliance.

Purpose

The **createbackupfs** command allocates space for queue manager back ups on the Appliance RAID volume. The storage is visible in the directory `mqbackup:///QMgrs`.

Syntax

```
▶▶ createbackupfs  — -s — size  ▶▶
```

Parameters

-s size

Specifies the size of the back up allocation in GB. You can specify a value in MB by entering the value followed by the character M. For example, to specify a size of 3 GB, enter 3. To specify a size of 1024 MB, enter 1024M.

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- You must run this command before you back up any queue managers.

Examples

- The following command allocates 4 GB of storage in the `mqbackup:///QMgrs` directory.

```
createbackupfs -s 4
```

Related commands

- [“deletebackupfs: clear previously allocated back up space” on page 666](#)
- [“mqbackup: back up queue manager” on page 667](#)
- [“mqrestore: restore queue manager from back up” on page 668](#)

deletebackupfs: clear previously allocated back up space

Free the space that was previously allocated for back up archive files on the appliance.

Purpose

The **deletebackupfs** command clears the storage on the appliance RAID volume previously allocated for back ups.

Syntax

```
▶▶ deletebackupfs  ▶▶
```

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 directory `mqbackup: ///QMGRS` must be empty. If it contains any files, the command will fail. You can use the delete command to delete files, see [“delete” on page 790](#).

Related commands

- [“createbackupfs:create appliance storage location for back up ” on page 666](#)
- [“mqbackup: back up queue manager” on page 667](#)
- [“mqrestore: restore queue manager from back up” on page 668](#)

mqbackup: back up queue manager

Back up a queue manager.

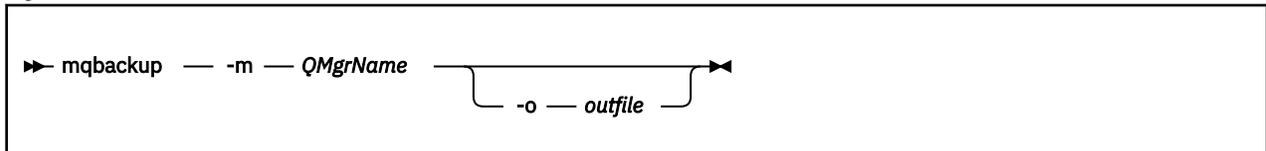
Purpose

You can use the **mqbackup** command to back up a queue manager, including all its log files and data. The command creates an archive and writes it to a location in the appliance file store. You must run the **createbackupfs** command to assign space for them, before you run the **mqbackup** command.

You can use the **mqbackup** when a queue manager is stopped, or when it is running. If you are backing up so that you can use an archive file to migrate the queue manager, or if you want to be able to restore a queue manager to the state it was in at a particular time, then you should stop the queue manager before you back it up. Taking a back up of a running queue manager requires more free disk space than backing up a stopped queue manager, see [“Usage notes” on page 667](#).

If the queue manager is running when you issue the **mqbackup** command, a warning message is displayed.

Syntax



Parameters

-m QMGrName

Specifies the name of the queue manager that you want to back up.

This parameter is required.

-o outfile

Optionally specifies the name of the back up file. If no name is specified, the filename `QMGrName.bak` is used. If the file already exists, no back up is made and an error is reported.

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.
- You can back up a queue manager while it is running, but this requires sufficient unallocated space on the disk to contain a temporary snapshot of the queue manager. This space is not required if the queue manager is stopped before the back up is taken.

- If a queue manager is stopped before taking a back up, it is locked for the duration of the back up and cannot be started, deleted or otherwise changed while the back up is running.
- The back up can take some time to run, during which period your CLI session will be suspended. Interrupting the CLI session will terminate the backup process.
- You can run the back up command on the primary instance of a queue manager on the main appliance in a disaster recovery configuration, or on the secondary instance on the recovery appliance. However, if synchronization is in progress when you try to back up a secondary instance, or if the data has become inconsistent on the secondary instance, the back up will fail.

Examples

- The following command backs up the queue manager QM1 to the file `safeandsound.bak`.

```
mqbackup -m QM1 -o safeandsound.bak
```

Related commands

- [“createbackupfs: create appliance storage location for back up”](#) on page 666
- [“deletebackupfs: clear previously allocated back up space”](#) on page 666
- [“mqrestore: restore queue manager from back up”](#) on page 668

mqrestore: restore queue manager from back up

Restore a queue manager from a back up archive.

Purpose

You can use the **mqrestore** command to restore a queue manager, including all its log files and data, from a previously taken back up. The command cannot run if there is already a queue manager with the same name on the appliance. The archive file must be located in the backupfs location on the appliance that was specified by the **createbackupfs** command.

Syntax

```
➤ mqrestore — -f — filename ➤
```

Parameters

-f filename

Specifies the name of the queue manager that you want to restore.

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- After restoration, the queue manager has the same name, configuration, and data as the original queue manager had when the archive was created. But any high availability or disaster recovery configuration is lost, the queue manager is restored as a stand-alone queue manager.
- Only one queue manager can be restored at a time.

Examples

- The following command restores the queue manager QM1 from the file `safeandsound.bak`.

```
mqrestore -f safeandsound.bak
```

Related commands

- [“createbackupfs:create appliance storage location for back up”](#) on page 666
- [“deletebackupfs: clear previously allocated back up space”](#) on page 666
- [“mqbackup: back up queue manager”](#) on page 667

High availability commands

crthagrp (create HA group)

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**.

Syntax

```
►► crthagrp — -s — SecretText — -a — IPAddress ◄◄
```

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.

-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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 `app11` and `app12` where a new, unique key is generated for communication between the appliances. The HA group primary interface of `app12` has the IP address `192.0.2.8`, the HA group primary interface of `app11` has the IP address `192.0.2.7`.

The following command is run from app11:

```
prepareha -s AGEW1823510HH -a 192.0.2.8
```

The following command is run from app12:

```
crthgrp -s AGEW1823510HH -a 192.0.2.7
```

Related reference

[“prepareha \(prepare appliance for HA group\)” on page 674](#)

Prepares an appliance to be part of an HA group that uses a unique, generated key for communication between appliances.

crthkeys (create HA keys)

Regenerate SSH public and private keys and exchange public keys between appliances in an HA pair.

Purpose

Good security practise requires that the secret keys used to secure communications between appliances in an HA pair are periodically regenerated and exchanged. Use the **crthkeys** command to regenerate and exchange the keys without affecting the existing HA queue managers.

Syntax

```
►► crthkeys ◄◄
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- You can run the **crthkeys** command at any time on either appliance in the HA pair, provided that no other HA command is running. The command does not have any operational or performance impact.

Related reference

[“dsphkeys \(display HA key information\)” on page 671](#)

Display information about the SSH keys used for secure communication between an HA pair.

dsphgrp (display HA group status)

Displays the status of the appliances in the high availability (HA) group.

Purpose

You can use the **dsphgrp** 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 **sethgrp -s** command.

V 9.1.0 If you are upgrading a Version 8.0 or Version 9.0 appliance to Version 9.1, the migration pending status is displayed until you run the `sethgrp -u` command to upgrade each HA queue manager, and the other appliance is upgraded to Version 9.1 too.

Syntax

```
▶▶ dsphagr ▶◀
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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:

```
dsphagr▶
```

Related reference

[“sethagr \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

dsphakeys (display HA key information)

Display information about the SSH keys used for secure communication between an HA pair.

Purpose

Good security practise requires that the secret keys used to secure communications between appliances in an HA pair are periodically regenerated and exchanged. Use the **dsphakeys** command to see when the keys used by the HA configuration were last generated.

Syntax

```
▶▶ dsphakeys ▶◀
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- If you run **dsphakeys** on an appliance on which HA was configured before the upgrade to Version 9.0.2, an error occurs. To correct the error, regenerate the keys by using the **crthakeys** command.
- You can run the **dsphakeys** command at any time on either appliance in the HA pair. The command does not have any operational or performance impact.
- The command displays the date and time that the keys were last generated in UTC time (UTC+00:00).

Related reference

[“crthakeys \(create HA keys\)” on page 670](#)

Regenerate SSH public and private keys and exchange public keys between appliances in an HA pair.

V 9.1.0 *dsphalink*

Display details of the link configured for HA replication.

Purpose

Displays the current settings for HA replication interface, and confirms that the IP address looks well configured.

Syntax

```
➤ dsphalink ➤
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Example

The following example shows the output of a **dsphalink** command:

```
Replication:          127.0.0.1 (10)
The dsphalink command succeeded.
```

dlthgrp (delete HA group)

Deletes a high availability (HA) group.

Purpose

You can use the **dlthgrp** command to delete an existing HA group.

Syntax

```
➤ dlthgrp ➤
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 267.
- 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 (make HA appliance primary)

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. The queue manager is started on the preferred appliance, if one is specified, with the 'winning' queue manager data set.

Note: When you run the **makehaprimary** command the data on the other appliance is discarded. See [“Resolving a partitioned problem in a high availability configuration” on page 561](#).

Syntax

```
▶▶ makehaprimary — QMname ▶▶
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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
```

Related concepts

[“Resolving a partitioned problem in a high availability configuration” on page 561](#)

A partitioned problem occurs when the two appliances in a high availability configuration lose the ability to communicate with each other. If both the primary and secondary connections are lost, the queue manager will run on both appliances at the same time.

prepareha (prepare appliance for HA group)

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

```
►► prepareha — -s — SecretText — -a — IPAddress — -t — timeout ◀◀
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 267.

Examples

- The following example shows the creation of an HA group for appliances `app11` and `app12` where a new, unique key is generated for communication between the appliances. The HA group primary interface of `app12` has the IP address `192.0.2.8`, the HA group primary interface of `app11` has the IP address `192.0.2.7`.

The following command is run from `app11`:

```
prepareha -s AGEW1823510HH -a 192.0.2.8
```

The following command is run from app12:

```
crthagrp -s AGEW1823510HH -a 192.0.2.7
```

Related reference

“[crthagrp \(create HA group\)](#)” on page 669
Creates a high availability (HA) group of appliances.

sethagr (set HA group)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

Purpose

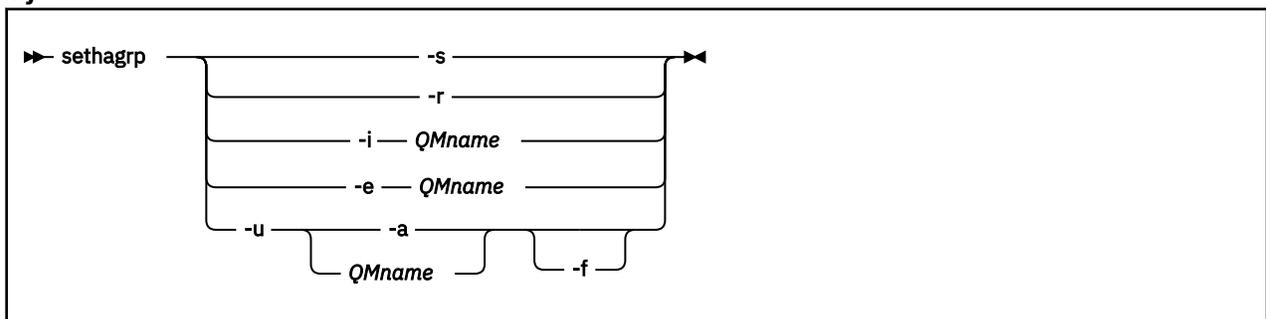
You use the **sethagr** command to pause the high availability (HA) group on an appliance. Any queue managers running on that appliance fail over to the other appliance in the group. You can then use **sethagr** to resume a previously paused HA group on the appliance.

You can also use the **sethagr** 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 stand-alone 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 **sethagr** command, see “[dltdrprimary](#)” on page 682.

V 9.1.0 When you upgrade an appliance from version 9.0.x to version 9.1, you must run **sethagr** to upgrade your HA queue managers.

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.

-u QMname [-f] | -u -a [-f]

Upgrade a queue manager from Version 8.0 or Version 9.0 to run in Version 9.1. You can upgrade single queue managers by specifying the queue manager name, or you can upgrade all HA queue managers by specifying the -a option (use -a with caution, it is recommended that you upgrade one queue manager at a time).

Specify the -f option to force the completion of the upgrade even if the current appliance cannot communicate with the other appliance in the HA group over either of the replication links.



CAUTION: Use the -f option only if you have an urgent need to run the queue manager, because you will likely have partitioned data when you restore the other appliance.

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 `mqc1i(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Related concepts

[“Upgrading a version 8.0 IBM MQ Appliance to version 9.1” on page 177](#)

You can upgrade a version 8.0 IBM MQ Appliance (running version 8.0 queue managers) to version 9.1 (running version 9.1 queue managers).

[“Upgrading a version 9.0 IBM MQ Appliance to version 9.1” on page 183](#)

You can upgrade a version 9.0 IBM MQ Appliance (running version 9.0 queue managers) to version 9.1 (running version 9.1 queue managers).

Related tasks

[“Adding an existing queue manager to a high availability group” on page 276](#)

You can add an existing queue manager to a high availability (HA) group by using the **sethagrp** command on the command line.

[“Removing a queue manager from a high availability group” on page 280](#)

You can remove a queue manager from a high availability (HA) group and run it as a stand-alone queue manager by using the **sethagrp** command on the command line.

[“Replacing a failed node in a high availability group” on page 286](#)

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.

Related reference

[“crthagrp \(create HA group\)” on page 669](#)

Creates a high availability (HA) group of appliances.

[“strmqm \(start queue manager\)” on page 610](#)

Start a queue manager.

[“dspmq \(display queue managers\)” on page 585](#)

Display information about queue managers.

[“status” on page 881](#)

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.

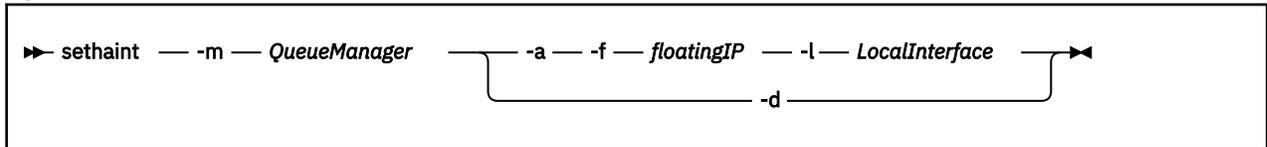
sethaint

Specify a floating IP address for a high availability (HA) queue manager, or delete an existing floating IP address.

Purpose

You use the **sethaint** command to specify a floating IP address that can be used by applications to connect to an HA queue manager, regardless of which appliance in an HA group it is actually running on. You also use **sethaint** to delete an existing floating IP address.

Syntax



Parameters

-m *QueueManager*

Identifies the HA queue manager that you are creating or deleting the floating IP address for.

-a

Specifies that you are adding the address specified by the -f and -l options.

-d

Specifies that you are deleting the floating IP address for the specified queue manager.

-f *floatingIP*

Specifies the floating IP address. You must specify the IP address using ip v4 dotted decimal notation (for example, "192.0.2.8").

-l *LocalInterface*

Specifies the name of the local interface that is used to connect to the queue manager on the two appliances in the HA group. For example, eth22. You can also specify an aggregate interface name, provided that the name obeys certain naming rules (see [“link-aggregation”](#) on page 803).

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(mqc.li)#`. To enter the IBM MQ administration mode, enter `mqc.li` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- You can run the **sethaint** command only once for each queue manager. You can only define one floating IP address for IBM MQ traffic on a queue manager
- Both appliances in the HA pair must be active when you run this command.
- The local interface that you specify can be a physical interface that exists on both appliances and has a static IP address configured. You can also specify an aggregate interface that exists with the same name on each appliance. The name of the aggregate interface must obey certain rules to be used for this purpose, see [“link-aggregation”](#) on page 803.
- The floating IP address must be a valid IPv4 address that is not already defined on either appliance, and it must belong to the same subnet as the static IP addresses defined for the local interface. See [“link-aggregation”](#) on page 803 for naming details.

Example

The following example shows the floating IP address 192.0.2.15 being allocated for queue manager QM1 and associated with the local interface eth22:

```
sethaint -m QM1 -a -f 192.0.2.15 - l eth22
```

► V 9.1.0 **sethalink**

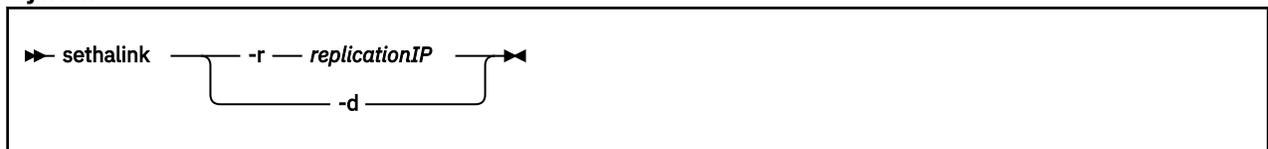
Specify a non-default link to be used for HA replication.

Purpose

By default, high availability configurations use the 10 Gb link eth21 for replicating data between the two appliances in the HA group.

Advanced users, who are networking experts, can choose to use a different interface for replication. The **sethalink** command is used to configure the alternative interface. ► M 2002 You might, for example, choose to use one of the 40 Gb links for replication.

Syntax



Parameters

-r replicationIP

Specifies the IP address of the interface to be used for replication. You must specify the IP address using ip v4 dotted decimal notation (for example, "192.0.2.8").

-d

Return to the default configuration.

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- The IP address specified must map to a valid local interface or aggregation interface, otherwise the command fails. (You cannot specify a VLAN interface.)
- The local interface should be at least a 10 Gb link, otherwise you are warned that performance might be poor.
- If an HA configuration is already active, the command is rejected and an error message issued.

Example

The following example shows the IP address 192.0.2.15 being allocated for replication in the HA configuration:

```
sethalink -r 192.0.2.15
```

sethapreferred (set HA preferred appliance)

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

```
▶▶ sethapreferred — QMname ▶▶
```

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.

Related reference

[“sethagr \(set HA group\)” on page 675](#)

V 9.1.0 Pauses and resumes a high availability group on an appliance. Removes or adds existing queue managers from or to a high availability group. Upgrades HA queue managers.

[“clearhapreferred \(clear HA preferred appliance\)” on page 679](#)

Clears the preferred appliance for a queue manager in a high availability (HA) group.

clearhapreferred (clear HA preferred appliance)

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

```
▶▶ clearhapreferred — QMname ▶▶
```


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 have not been 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.

-f floatingIP

If you are configuring disaster recovery for a high availability pair, you can optionally specify a floating IP address. The queue manager specified by `-m QMName` must already belong to an HA pair if you use the `-f` option. 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.

You do not have to 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.

Used with this option, the `crtdrprimary` command configures the queue manager on both appliances in the HA pair, and reserves storage for the data snapshot on both appliances.

Note: For versions earlier than IBM MQ Appliance Version 9.1.5 the specification of the floating IP address is mandatory for an HA pair, and appliances in the HA group must be in the same subnet.

In Version 9.1.5, you can currently configure only one HA/DR queue manager with no floating IP address, contact IBM Support for help with overcoming this limitation.

Usage Notes

- The queue manager must be stopped before you run `crtdrprimary`. You can use the `endmqm` command to stop the queue manager.
- On successful completion, the command outputs the `crtdrsecondary` command that you must run on the recovery appliance to configure the queue manager on there.
- There must be sufficient memory for both the queue manager data and the snapshot of the data that is required for disaster recovery. For example, if there is 100 GB free and you create a queue manager with the default size of 64 GB, and then run the `crtdrprimary` command for that queue manager without freeing some additional space, `crtdrprimary` fails because there is not another 64 GB of free space.
- After the command completes, you can restart the queue manager on the live appliance. This can continue to run while messages are replicated to the recovery appliance.

Examples

The following example shows the existing queue manager QM1 being prepared for running on a disaster recovery configuration, with the appliance that you run the command on as the live system, and the appliance named `mydrapp1` as the recovery.

```
crtdrprimary -m QM1 -r mydrapp1 -i 198.51.100.0 -p 2015
```

Upon successful completion, the command returns the following information:

```
Queue manager QM1 is prepared for Disaster Recovery replication.  
Now execute the following command on appliance mydrapp1:  
crtdrsecondary -m QM1 -s 65536 -l myliveapp1 -i 198.51.100.24 -p 2015
```

The following example shows the high availability queue manager QM3 being prepared for running on a disaster recovery configuration, with the appliance that you run the command on as the live system, and the appliance named `myliveapp3` as the recovery. In this example the eth20 port on the HA appliance

currently running QM3 has the static IP address 198.51.100.20 (which is not used in the command) and the floating IP address 198.51.100.10. The DR appliance has the IP address 198.51.100.124.

```
crtdrprimary -m QM3 -r mydrapp3 -i 198.51.100.124 -p 2015 -f 198.51.100.10
```

Upon successful completion, the command returns the following information:

```
Queue manager QM3 is prepared for Disaster Recovery replication.  
Now execute the following command on appliance mydrappl:  
crtldrsecondary -m QM3 -s 65536 -l myliveapp3 -i 198.51.100.10 -p 2015
```

Related reference

[“crtdrsecondary” on page 682](#)

Creates a secondary version of a queue manager on the recovery appliance in a disaster recovery configuration.

crtdrsecondary

Creates a secondary version of a queue manager on the recovery appliance in a disaster recovery configuration.

Purpose

All parameters are supplied by the equivalent **crtdrprimary** command and should be entered exactly as shown in the output from that command.

After this command is run, synchronization of data from the main to the recovery appliance begins. The queue manager status is shown as stopped, and initial synchronization progress can be followed by using the status command.

Related reference

[“crtdrprimary” on page 680](#)

Augments an existing queue manager to become the primary queue manager in a disaster recovery configuration.

[“status” on page 881](#)

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.

dltdrprimary

Remove a queue manager currently in the primary role from DR control.

Purpose

You use the **dltdrprimary** command to remove a queue manager from the disaster recovery configuration on the appliance. The queue manager is in the primary role. The queue manager must have the stopped status on the appliance. You receive an error if you run the command on a queue manager that is running, or a queue manager that is in the secondary role.

After you run **dltdrprimary**, the queue manager is left as a stand-alone queue manager and can be started or deleted as required.

If you run **dltdrprimary** on a queue manager that is running in a high availability group, the command removes the disaster recover status and the space reserved for the snapshot logical volume on both appliances in the HA group.

Syntax

```
►► dltdrprimary — -m — QMName ◄◄
```


Example

The following example shows the output of a **dspdrlink** command with a queue manager parameter:

```
mqa(mqcli)# dspdrlink -m QM1
The replication link for queue manager 'QM1' is IP address
(172.20.4.0) interface (eth20)
The dspdrlink command succeeded.
```

The following example shows the output of a **dspdrlink** command with no parameters:

```
mqa(mqcli)# dspdrlink
The default replication link is IP address (172.20.4.0) interface
(eth20)
QMGR  DR Interface IP (Name)
-----
QM1   172.20.2.0 (eth21)
QM2   172.20.3.0 (eth22)
The dspdrlink command succeeded.
```

makedrprimary

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

Purpose

You use the **makedrprimary** command on an appliance to identify it as the primary version in a disaster recovery configuration.

You should always check the disaster recovery status of a queue manager before you issue the **makedrprimary** command on that queue manager.

If you run **makedrprimary** when the queue manager is in the partitioned state (that is, each appliance has a different version of the queue manager data) the version of the queue manager and associated data on this appliance are identified as the definitive version.

If you run **makedrprimary** on the recovery appliance when the secondary queue manager is inconsistent (that is, replication has not completed successfully and the queue manager would be unable to start), then the command starts the process of reverting the queue manager to the data snapshot taken before the queue manager became inconsistent. You can monitor the progress of the reversion by using the `status` command, see [“status” on page 881](#). If the reversion is interrupted for any reason, it will resume and complete. After it has reverted to the snapshot, the queue manager becomes the primary version in the disaster recovery configuration.

If you run **makedrprimary** on a secondary queue manager after initial synchronization has failed, a message informs you that you cannot do this until the initial synchronization has completed. If the main appliance has failed and will not be restored (so that the initial synchronization can never complete), then the queue manager must be deleted on the recovery appliance by running the **dltdrsecondary** command.

Syntax

```
► makedrprimary — -m — QMName ◄
```

Parameters

-m QMName

Specifies the queue manager that you are identifying as the primary queue manager in a disaster recovery configuration.

Related reference

[“makedrsecondary” on page 685](#)

Prevents a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.

makedrsecondary

Prevents a queue manager on an appliance in a disaster recovery configuration from starting, and specifies that it has the secondary role.

Purpose

You use the **makedrsecondary** command on an appliance to prevent a queue manager on that appliance from starting. If you attempt to start the queue manager by using the **strmqm** command, you receive an error message. You identify the queue manager as the secondary version in a disaster recovery configuration. If you run this command when the queue manager is in the partitioned state, data on this appliance associated with the queue manager is discarded. When the **makedrprimary** command is run on the other appliance in the disaster recovery configuration, the queue managers are resynchronized and data is replicated from the primary queue manager to the secondary.

Syntax

```
➔ makedrsecondary -m QMName ➔
```

Parameters

-m QMName

Specifies the queue manager that you are identifying as the secondary queue manager in a disaster recovery configuration.

Related reference

[“makedrprimary” on page 684](#)

Switches a disaster recovery queue manager to have the primary role in the disaster recovery configuration.

V 9.1.0 setdrlink

Specify a non-default link to be used for DR replication in an HA configuration.

Purpose

By default, high availability configurations use the 10 Gb link eth20 for replicating data between the HA group and a DR appliance.

Advanced users, who are networking experts, can choose to use a different interface for replication. The **setdrlink** command is used to configure the alternative interface. **M 2002** You might, for example, choose to use one of the 40 Gb links for replication.

The command sets a new interface from that moment forward, so that any **crtdrprimary** and **crtdrsecondary** commands set up a DR configuration using the newly specified interface. Any queue managers configured to use DR before this point will continue to use the default interface, rather than the new one, unless you remove and recreate the DR configuration for those queue managers.

Syntax

```
➔ setdrlink -r replicationIP -d ➔
```

Parameters

-r replicationIP

Specifies the IP address of the interface to be used for replication. You must specify the IP address using ip v4 dotted decimal notation (for example, "192.0.2.8").

-d

Return to the default configuration.

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 IP address specified must map to a valid local interface or aggregation interface, otherwise the command fails. (You cannot specify a VLAN interface.)
- The local interface should be at least a 10 Gb link, otherwise you are warned that performance might be poor.
- You must select the same interface for both appliances in your HA group.

Example

The following example shows the IP address 192.0.2.15 being allocated for replication in the HA configuration:

```
setdrlink -r 192.0.2.15
```

Troubleshooting commands

dltmqras (*delete mq ras*)

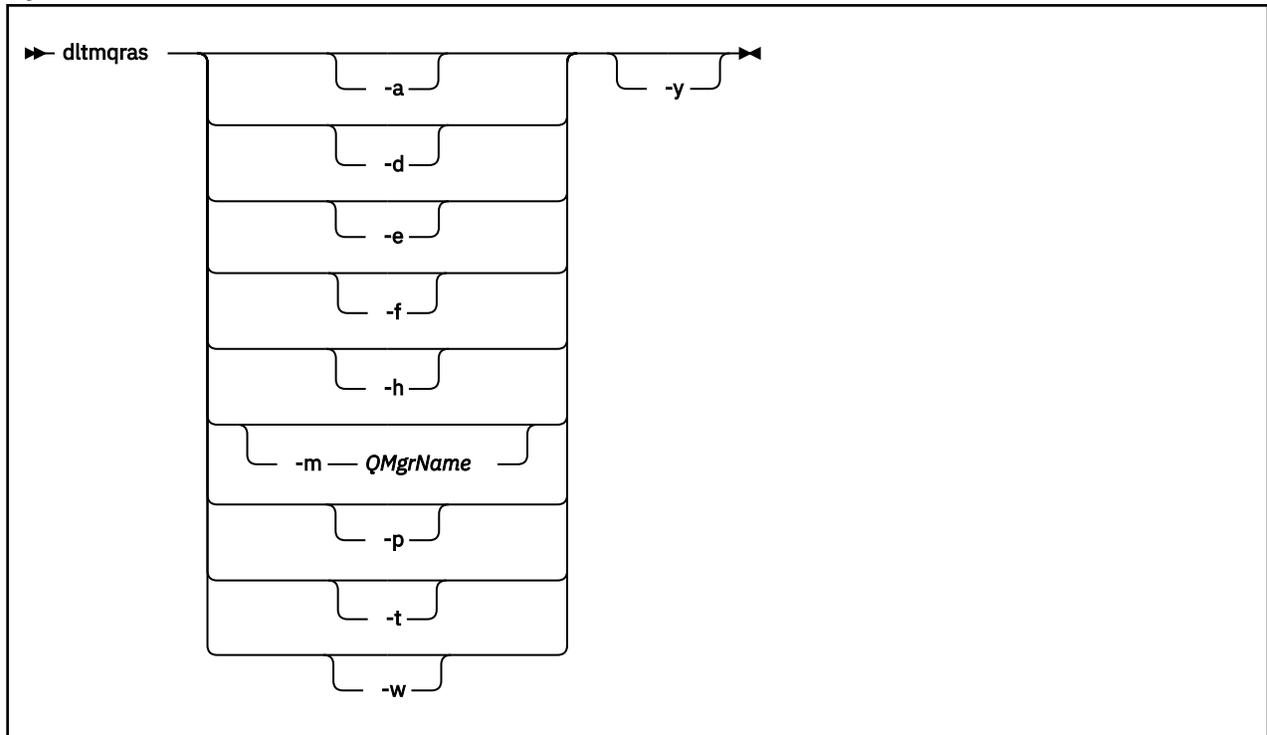
Deletes the IBM MQ error log files.

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 (display mq error log)

Displays the IBM MQ error log files.

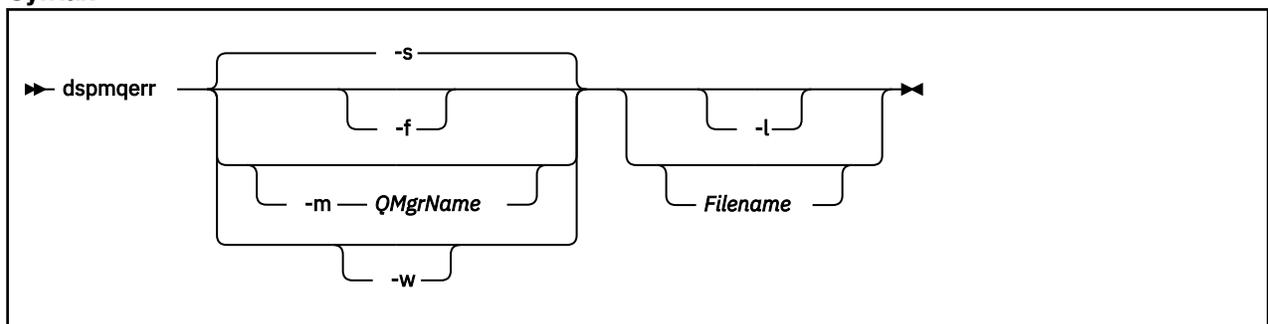
Purpose

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:

```
dspmqrtr  
dspmqrtr -s  
dspmqrtr MQSystem.log
```

- The following command lists all the error logs on the appliance (but not FDC files):

```
dspmqrtr -l
```

- The following command lists all the FDC files:

```
dspmqrtr -f -l
```

- The following command lists all the IBM MQ Console files:

```
dspmqrtr -w -l
```

- The following command lists all the log files for the queue manager QM1:

```
dspmqrtr -m QM1 -l
```

- The following command lists the first log file in the log file directory for the queue manager QM1:

```
dspmqrtr -m QM1
```

- The following command lists the log file for the queue manager QM1 named AMQERR02.LOG:

```
dspmqrtr -m QM1 AMQERR02.LOG
```

- The following command displays the FDC file named AMQ12345.FDC:

```
dspmqrtr -f AMQ12345.FDC
```

- The following command displays the IBM MQ Console file named messages.log:

```
dspmqrtr -w messages.log
```

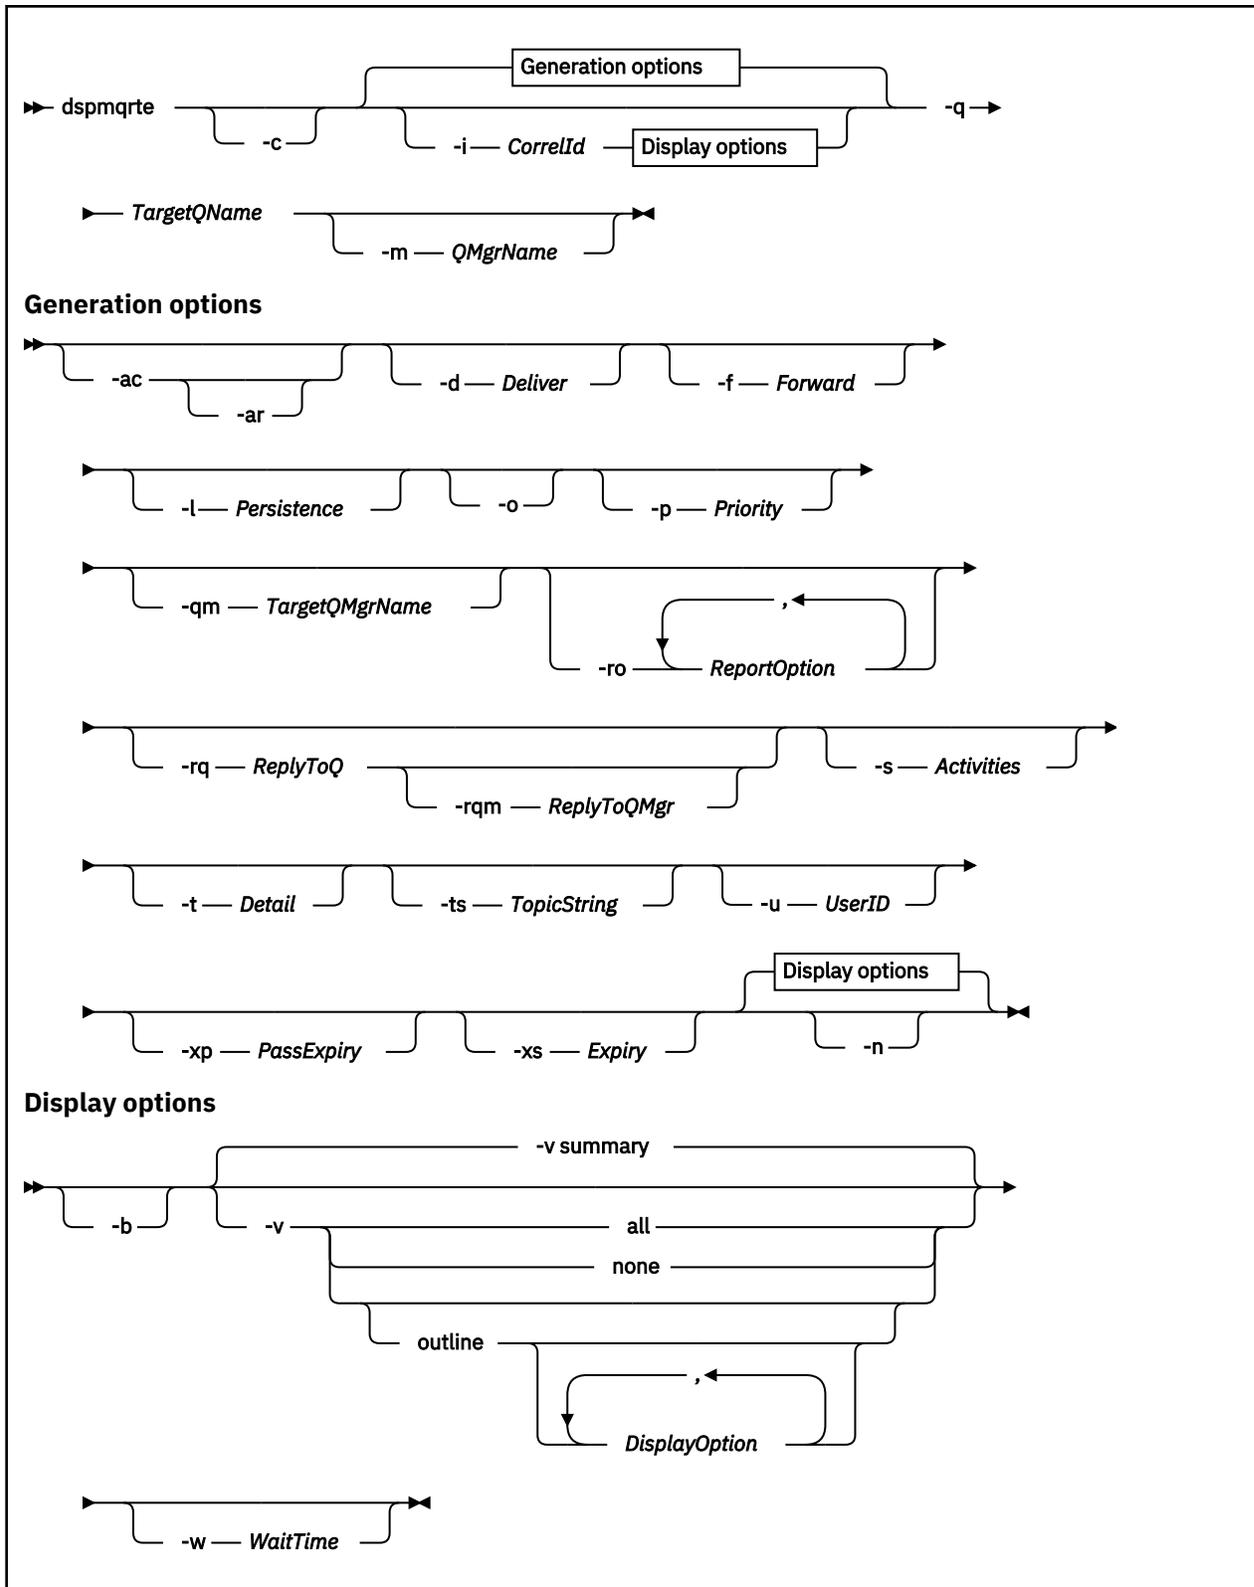
dspmqrte (display route information)

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



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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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.

V 9.1.5 *dspmqrtrc* (display formatted trace)

Format and display IBM MQ trace.

Purpose

Use the **dspmqrtrc** command to display IBM MQ formatted trace output.

Syntax



Required parameters

InputFileName

The name of the file containing the unformatted trace, for example:

```
AMQ12345.01.TRC
```

The output file is named AMQ *yyyyy.zz.FMT*, based on the PID of the trace file and is written to the `mqtrace://` directory.

Specify an asterisk character (*) to format all trace files.

Optional parameters

-h

Omit header information from the report.

-s

Extract trace header and put to stdout.

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.
- Trace files are created in the `mqtrace://` directory.
- Formatted files are written to the `mqtrace://` directory.

Related commands

Table 52. Related command names and descriptions

Command	Description
“endmqtrc (end trace)” on page 599	End trace
“strmqtrc (Start trace)” on page 613	Start trace

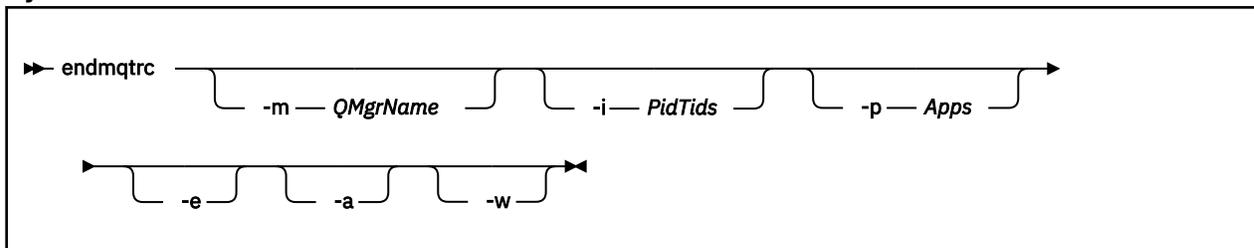
endmqtrc (end trace)

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 \(Start trace\)” on page 613](#)

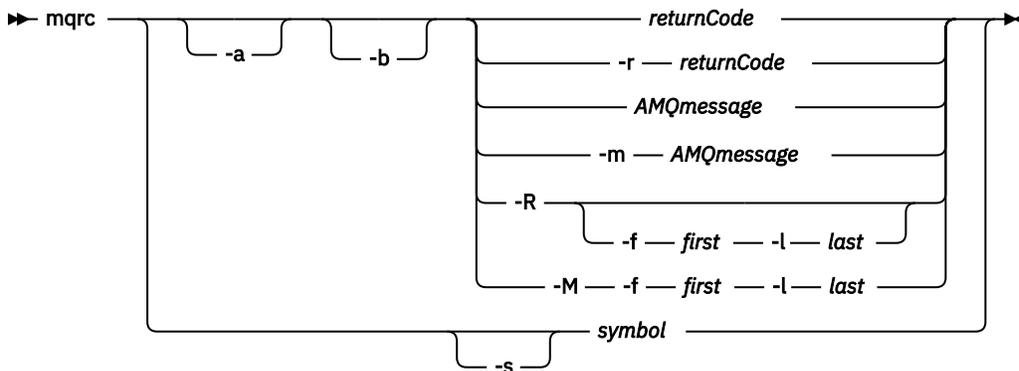
mqrc

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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

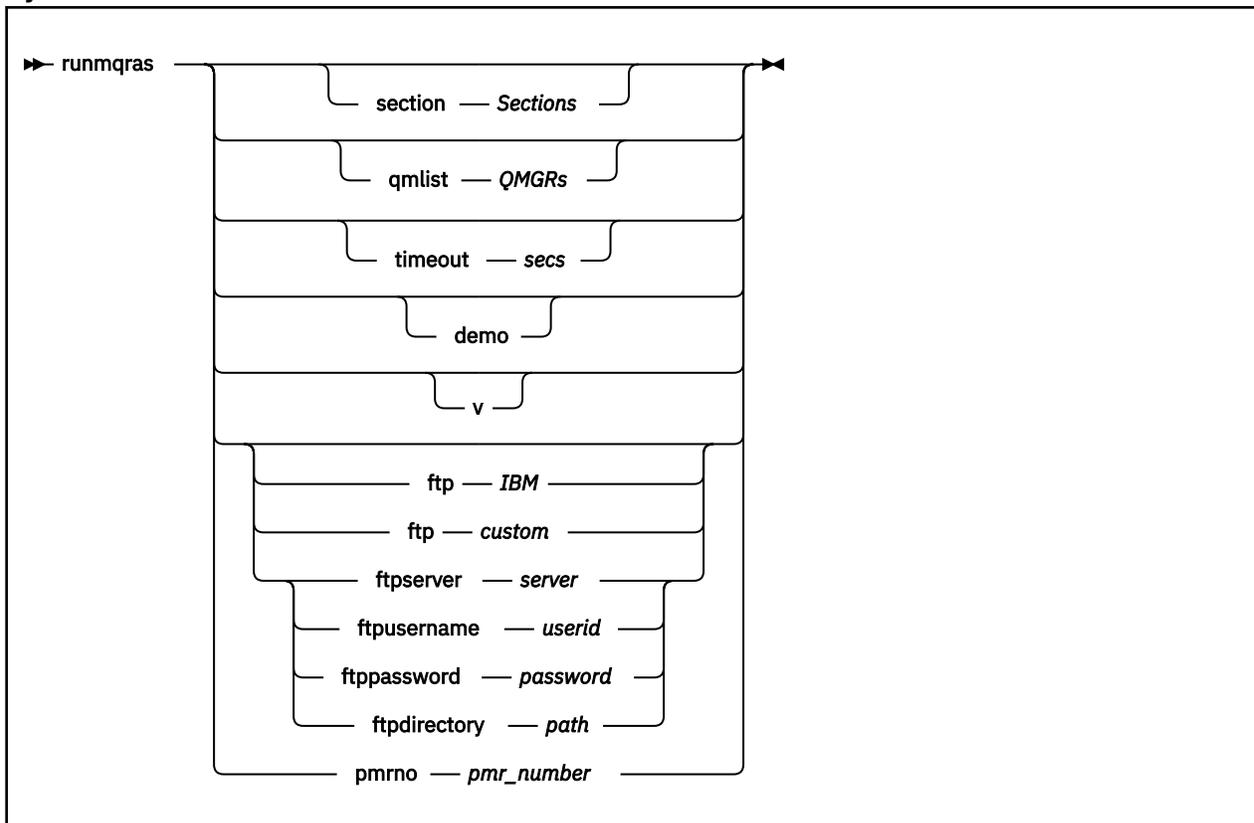
You can use the **runmqras** command to gather diagnostic information from the appliance into a single archive. You can use this command to gather information about an application or appliance failure, possibly for submission into IBM when you report a problem.

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 788), or by using the IBM MQ Appliance web UI (see “Managing files by using the IBM MQ Appliance web UI” on page 402). 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.

You could also consider using the `qmlist` parameter with this option to limit the output to specified queue managers.

cluster

Gather information specific for clustering

defs

Gather the queue manager definitions and status

format

Generate and include formatted trace files.

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 `mqc(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` on the command line. To exit the IBM MQ administration mode, enter `exit` on the command line.
- Use **V9.1.5 runmqras -section format** to generate and include formatted trace files.
- 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)

Start trace at a specified level of detail, or report the level of tracing in effect.

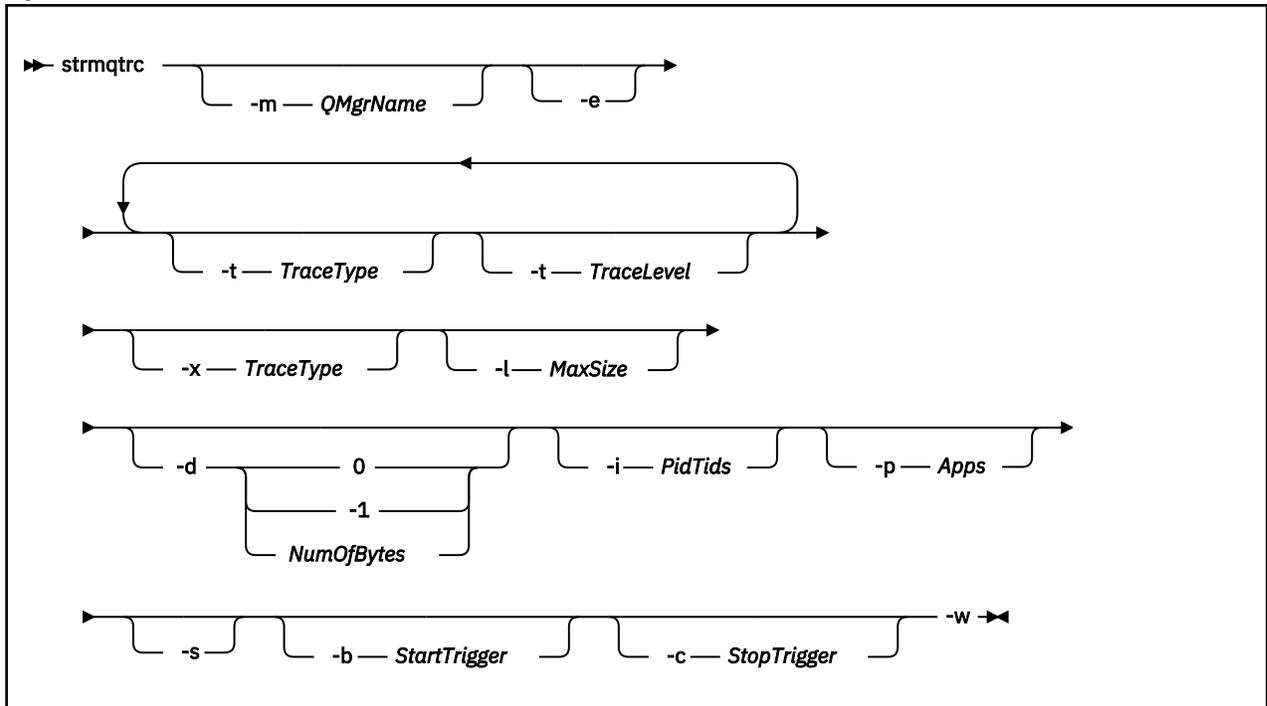
Purpose

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 788](#)), or by using the IBM MQ Appliance web UI (see [“Managing files by using the IBM MQ Appliance web UI” on page 402](#))

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 `a11`.

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(mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 **enmqmqtrc** command, are interpreted as having a logical OR between them. The maximum number of concurrent **strmqmqtrc** commands that can be in effect at one time is 16.
- When a trace file reaches the specified maximum, it is renamed to `AMQppppp.qq.TRS` and a new `AMQppppp.qq.TRC` file is started. If a previous copy of an `AMQppppp.qq.TRS` file exists, it is deleted.
- For more information about this command in IBM MQ, see [strmqmqtrc](#) 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.

```
strmqmqtrc -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.

```
strmqmqtrc -m exampleQM -x ssl -t parms
```

- The following command enables high-detail tracing of the processing flow for all components:

```
strmqmqtrc -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`:

```
strmqmqtrc -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 enmqmqtrc` 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 \(end trace\)” on page 599](#)

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
“clear intrusion-detected” on page 709	Clears the detected intrusion when the appliance is in Fail-Safe mode.
“clock” on page 710	Sets the date or time.
“configure terminal” on page 710	Enters configuration mode.
“diagnostics” on page 711	Enters Diagnostics mode. Use this command only at the explicit direction of IBM Support.
“disconnect” on page 711	Closes the session for an active user.
“echo” on page 711	Echoes text to the console.
“exec” on page 712	Calls and runs a configuration script.
“exit” on page 712	Applies configuration changes to the running configuration and returns to the parent mode.

Command	Description
“help” on page 713	Displays online help.
“login” on page 713	Logs a specific account on to the appliance.
ntp	Identifies an NTP server.
“ping” on page 714	Determines whether the network can reach a remote target.
“show” on page 715	Displays configuration or status information
“shutdown” on page 715	Restarts or shuts down the appliance.
“test tcp-connection” on page 717	Tests the TCP connection to a remote target.

The following commands are available in most configuration modes:

Command	Description
“admin-state” on page 708	Sets the administrative state for the configuration.
“cancel” on page 709	Exits this configuration mode without saving changes to the running configuration.
“disconnect” on page 711	Closes the session for an active user.
“echo” on page 711	Echoes text to the console
“exit” on page 712	Applies configuration changes to the running configuration and returns to the parent mode.
“help” on page 713	Displays online help
“ping” on page 714	Determines whether the network can reach a remote target.
“reset” on page 714	Restores the default values to the configuration.
“show” on page 715	Displays configuration or status information
“test tcp-connection” on page 717	Tests the TCP connection to a remote target.
“traceroute” on page 718	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
```

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
```

Related reference

[“shutdown” on page 715](#)

This command restarts or shuts down the appliance.

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.

Related concepts

[“Global configuration commands” on page 792](#)

You can use the global configuration commands to manage configuring the IBM MQ Appliance.

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.

You must be logged in as the user `admin` to use this command.



Attention: Use this command only at the explicit direction of IBM Support.

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 ...
...
# disconnect 36
Session 36 closed
```

Related reference

[“show users” on page 911](#)

This command lists all users who are currently logged in to the appliance.

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.

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.

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.

Related reference

[“ip-preference” on page 765](#)

This command sets the preferred IP version when the DNS provider publishes both versions of addresses.

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
#
```

Related reference

[“exit” on page 712](#)

This command applies configuration changes to the running configuration and returns to the parent mode.

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.

Related concepts

[“Monitoring and reporting commands” on page 881](#)

You use the monitoring and reporting commands to view the status of queue manager and of the appliance.

shutdown

This command restarts or shuts down the appliance.

Syntax

shutdown

shutdown reboot [*seconds*]

shutdown halt [*seconds*] (deprecated)

shutdown poweroff [*seconds*]

Parameters

reboot

Shuts down and restarts the appliance.

halt

Shuts down the appliance without restarting. The power to the appliance remains on. This keyword is deprecated. Use **poweroff** instead.

poweroff

Stops the appliance and turns off the power.

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 after waiting ten seconds.

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 and turn off the appliance.

```
# shutdown poweroff 60
Shutdown in 60 second(s).
#
```

Related reference

[“boot config” on page 824](#)

This command designates the startup configuration for the next restart.

summary

This command specifies the descriptive summary for the configuration.

Syntax

summary "*string*"

Parameters

string

Specifies the descriptive summary.

Guidelines

The **summary** command specifies the descriptive summary for the configuration. Enclose the summary in double quotation marks.

Example

Add the summary to a configuration.

```
# summary "Amended server list"
```

template

This command runs an interactive command-line script.

Syntax

template *URL*

Parameters

URL

Specifies the fully qualified location of the interactive command-line script.

Guidelines

The **template** command specifies the URL of the interactive command-line script. The script is an XML file that can be local or remote to the appliance. The script must conform to the `store:///schemas/dp-cli-template.xsd` schema.

Example

Run the interactive script as defined in the `local:///shell-script.xml` file.

```
# template local:///shell-script.xml
#
```

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]
```

 **V 9.1.4**

Test the connection specifying the local IP address to use.

```
test tcp-connection address port [seconds][local-address]
```

or:

```
test tcp-connection host port [seconds] [IP-version] [local-address]
```

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.

V 9.1.4 local-address

Specifies the IP address of a local interface to use for the connection test.

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
#
```

- **V 9.1.4**

The following command verifies connectivity to port 1234 at the remote server myhost123. In this scenario, the routing table selects a network interface to connect to the remote server, then it sends the request from that interface's primary IP address.

```
mqa(config)# test tcp-connection myhost123 1234
```

The following command repeats the same test using the source IP address 9.20.49.83, which can be either a primary or secondary IP address associated with a local network interface.

```
mqa(config)# test tcp-connection myhost123 1234 9.20.49.83
```

Related reference

[“ip-preference” on page 765](#)

This command sets the preferred IP version when the DNS provider publishes both versions of addresses.

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.

You must be logged in as the admin user to use this command.

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
```

Related reference

[“ip-preference” on page 765](#)

This command sets the preferred IP version when the DNS provider publishes both versions of addresses.

[“ping” on page 714](#)

This command determines whether the network can reach a remote target.

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.

Related reference

[“group” on page 720](#)

This command associates an account with a user group.

[“password” on page 721](#)

This command manages the password for an account.

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.

Related reference

[“access-level” on page 720](#)

This command assigns an account type of an account.

[“group” on page 720](#)

This command associates an account with a user group.

snmp-cred

This command adds SNMP V3 credentials to the account.

Syntax

snmp-cred *engine-ID authentication-protocol authentication-secret-type authentication-secret privacy-protocol privacy-secret-type privacy-secret*

Parameters

engine-ID

Specifies the engine ID of the SNMP V3 engine for which this account is being defined. A value of 0 is the shorthand representation of the engine ID of the local SNMP V3 engine on the appliance. For any other engine ID, the value is a hex string that represents the 5 - 32-byte value.

authentication-protocol

Identifies which authentication protocol to use. The default value is sha.

none

The account has no authentication key.

md5

The account uses HMAC-MD5-96 as the authentication protocol.

sha

The account uses HMAC-SHA-96 as the authentication protocol.

authentication-secret-type

Indicates whether the authentication secret is a password or a fully localized key. This parameter is required when the authentication protocol is md5 or sha. The default value is password.

password

The authentication secret is a password that is converted to an intermediate key with a standardized algorithm, and then localized against the engine ID value.

key

The authentication secret is a fully localized key. Specifying a fully localized key is useful when the key was initially created on another system.

authentication-secret

Specifies the secret, or key, for authentication for this account. This parameter is required when the authentication protocol is md5 or sha.

- If a password, specify a plaintext password that is at least 8 characters long.
- If a key and HMAC-MD5 are the authentication protocol, specify the hex representation of a 16-byte key.
- If a key and HMAC-SHA-96 are the authentication protocol, specify the hex representation of a 20-byte key.

You can use colons (:) between every 2 hex characters.

privacy-protocol

Identifies which privacy (encryption) protocol to use. The default value is des.

none

The account has no privacy key.

des

The account uses CBC-DES as the privacy protocol.

aes

The account uses CFB128-AES-128 as the privacy protocol.

privacy-secret-type

Indicates whether the privacy secret is a password or a fully localized key. This parameter is required when the privacy protocol is des or aes.

password

The privacy secret is a password that is converted to an intermediate key with a standardized algorithm, and then localized against the engine ID value.

key

The privacy secret is a fully localized key. Specifying a fully localized key is useful when the key was initially created on another system.

privacy-secret

Specifies the secret, or key, for privacy (encryption) for this account. This parameter is required when the privacy protocol is des or aes.

- If a password, specify a plaintext password that is at least 8 characters long.
- If a key and HMAC-MD5 are the authentication protocol, specify the hex representation of a 16-byte key.
- If a key and HMAC-SHA-96 are the authentication protocol, specify the hex representation of a 20-byte key.

You can use colons (:) between every 2 hex characters.

Guidelines

The **snmp-cred** command adds SNMP V3 credentials for this account. Each account can have multiple SNMP V3 credentials, one for each SNMP V3 engine that is identified by an *engine-ID* parameter.

Note: The current implementation supports an SNMP V3 credential for the local engine ID only. Therefore, there can be only one SNMP V3 credential for each account.

The secret for authentication and privacy can be defined either as a password (passphrase) or as a localized hex key. If a password, it is hashed and localized with the engine ID.

Examples

- Create SNMP V3 credentials for this account on the appliance with HMAC-MD5-96 as the authentication algorithm, and DES-CBC as the privacy algorithm. The password aBigSecret is converted to a localized authentication key, and the password aDifferentSecret is converted to a localized encryption key.

```
snmp-cred 0 md5 password aBigSecret des password aDifferentSecret
```

- Create SNMP V3 credentials for this account on the remote machine with the engine ID 000000000000000000000002, with HMAC-MD5-96 as the authentication algorithm, and with no privacy algorithm. The password is maplesyrup, which is converted to a localized key for the specified engine ID (000000000000000000000002).

```
snmp-cred 000000000000000000000002 md5 password maplesyrup none password ""
```

- Create SNMP V3 credentials for this account on the remote machine with the engine ID 000000000000000000000002, with HMAC-MD5-96 as the authentication algorithm, and with no privacy algorithm. The fully localized key is 52:6f:5e:ed:9f:cc:e2:6f:89:64:c2:93:07:87:d8:2b.

```
snmp-cred 000000000000000000000002 md5 key  
52:6f:5e:ed:9f:cc:e2:6f:89:64:c2:93:07:87:d8:2b none password ""
```

suppress-password-change

This command control whether the password for this account must be changed after the initial login by the account owner.

Syntax

Account owner does not need to change the account passwords after initial login.

suppress-password-change on

Forces account owner to change the account passwords after initial login.

suppress-password-changeoff

Parameters

on

Indicates that the account owner does not need to change the account passwords after initial login.

off

Forces the account owner must change the account passwords after initial login. This setting is the default behavior.

Guidelines

The **suppress-password-change** commands control whether the password for this account must be changed after the initial login by the account owner. By default, all local users must change their passwords after initial creation.

Note: The property is available during only initial creation and is unavailable when you edit this configuration.

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
#
```

Audit Log Settings commands

Audit Log Settings commands set the size and the number of rotations of the audit log.

To enter the mode, use the Global `audit-log-settings` command.

Related reference

[“audit-log-settings” on page 793](#)

This command enters Audit Log Settings mode.

rotate

This command sets the number of rotations of audit log files.

Syntax

rotate *rotations*

Parameters

rotations

Sets the number of rotations of audit log files. Enter a value in the range 1 - 100. The default value is 3.

Guidelines

The **rotate** command sets the number of rotations of audit log files.

- When the contents of the `audit-log` file reach the size set by the **size** command, a rotation occurs. A new audit file continues to record audit events. The `audit-log` file that is filled becomes the `audit-`

log.1 file. More rotation files are renamed from audit-log.*n* to audit-log.*n*+1, for as many rotations set by the **rotate** command.

- When the maximum number of rotations are generated, the oldest rotation file is replaced. The data in the oldest rotation is lost.

Example

Set the appliance for 5 rotations of audit log files at 5000 KB each. In this case, the appliance can maintain up to 30,000 KB of audit records, which are the audit-log file and its five rotations.

```
# size 5000
# rotation 5
```

Related reference

[“size” on page 726](#)

This command sets the size of audit log files.

size

This command sets the size of audit log files.

Syntax

size *KB*

Parameters

KB

Sets the size for audit log files in KB. Enter a value in the range 250 - 500000. The default is 1000.

Guidelines

The **size** command sets the size of audit log files.

- When the contents of the audit-log file reach the size set by the **size** command, a rotation occurs. A new audit file continues to record audit events. The audit-log file that is filled becomes the audit-log.1 file. More rotation files are renamed from audit-log.*n* to audit-log.*n*+1, for as many rotations set by the **rotate** command.
- When the maximum number of rotations are generated, the oldest rotation file is replaced. The data in the oldest rotation is lost.

Example

Set the appliance for 5 rotations of audit log files at 5000 KB each. In this case, the appliance can maintain up to 30,000 KB of audit records, which are the audit-log file and its five rotations.

```
# size 5000
# rotation 5
```

Related reference

[“rotate” on page 725](#)

This command sets the number of rotations of audit log files.

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, complete the following steps:

1. From the appliance command line, enter global configuration mode:

```
config
```

2. From global configuration mode, enter crypto configuration mode:

```
crypto
```

3. Type `exit` to save your changes and leave crypto configuration mode, then type `exit` again to leave global configuration mode.

Note: For certificate commands for queue managers, see [“Queue manager security management commands”](#) on page 629

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.

Related concepts

[“Crypto commands”](#) on page 726

You can use the crypto commands to manage certificates on the IBM MQ Appliance.

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
```

Related reference

[“idcred” on page 733](#)

This command creates identification credentials.

[“key” on page 734](#)

This command creates an alias for a private key.

[“password-map” on page 740](#)

This command manages the encrypted passwords to a password aliases in a password map file.

[“valcred” on page 749](#)

This command enters Validation Credentials mode.

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.

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.

Syntax

Exports a certificate

crypto-export *cert name* *output file*

Parameters

cert name

Identifies the name of the certificate.

output file

Identifies the name and location to store the export.

Guidelines

The **crypto-export** command creates an export package that contains a certificate.

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-import

This command imports an export package that contains a certificate.

Syntax

Imports a certificate

crypto-import *cert alias* *input file*

Parameters

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.

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
```

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-140-2-11

Runs the firmware in FIPS 140-2 Level 1 mode.

Guidelines

The **crypto-mode-set** command sets the cryptographic mode for the appliance. This setting affects only the encryption used for system management aspects of the appliance. For example, the encryption of administrative user passwords, CLI, and Web UI connections to your MQ Appliance. It does not affect the encryption used for your MQ channel traffic, which is configured on a per queue manager basis.

Changes made using **crypto-mode-set** take 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.

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-140-2-11
```

- Change the cryptographic mode back to permissive mode at the next reload of the firmware.

```
# crypto-mode-set permissive
```

Related concepts

[“FIPS compliance” on page 494](#)

Gives a guide to FIPS 140-2 level 1 compliance on the IBM MQ Appliance.

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
```

Related reference

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

[“key” on page 734](#)

This command creates an alias for a private key.

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
```

Related reference

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

[“idcred” on page 733](#)

This command creates identification credentials.

[“password-map” on page 740](#)

This command manages the encrypted passwords to a password aliases in a password map file.

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 [key-type type] rsa {1024 | 2048 | 4096} [digest algorithm] [ecdsa-curve name] [gen-object] [object-name name] [gen-sscert] [days number-days] [file-name name] [config-file 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.

key-type type

Indicates the type of key to generate. The following values are supported. The default value is RSA.

RSA

Generates RSA keys.

ECDSA

Generates Elliptic Curve Digital Signature Algorithm (ECDSA) keys.

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. This keyword is required only when the key type set by the key-type keyword is RSA.

digest algorithm

Specifies the hash algorithm of the generated RSA keys. The default value is sha256. This keyword is required only when the key type set by the key-type keyword is RSA and the gen-sscert keyword is set to on.

sha1

Uses SHA-1.

sha256

Uses SHA-256.

ecdsa-curve name

Specifies the elliptic curve to use to generate the ECDSA keys. The curves in the following table are supported. The default value is secp256r1. This keyword is required only when the key type is ECDSA.

Curve name	Value	Reference
sect163k1	1	RFC 4492
sect163r1	2	RFC 4492
sect163r2	3	RFC 4492
sect193r1	4	RFC 4492
sect193r2	5	RFC 4492
sect233k1	6	RFC 4492
sect233r1	7	RFC 4492
sect239k1	8	RFC 4492
sect283k1	9	RFC 4492
sect283r1	10	RFC 4492
sect409k1	11	RFC 4492
sect409r1	12	RFC 4492
sect571k1	13	RFC 4492
sect571r1	14	RFC 4492
secp160k1	15	RFC 4492
secp160r1	16	RFC 4492
secp160r2	17	RFC 4492
secp192k1	18	RFC 4492
secp192r1	19	RFC 4492

<i>Table 53. Supported elliptic curves (continued)</i>		
Curve name	Value	Reference
secp224k1	20	RFC 4492
secp224r1	21	RFC 4492
secp256k1	22	RFC 4492
secp256r1	23	RFC 4492
secp384r1	24	RFC 4492
secp521r1	25	RFC 4492
brainpoolP256r1	26	RFC 7027
brainpoolP384r1	27	RFC 7027
brainpoolP512r1	28	RFC 7027

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.

config-file name

Optionally, when generating a CSR or self-signed certificate, specifies the configuration file to read the DN and OIDs from. The configuration file must be in the OpenSSL CNF format. You must upload the configuration file to the `temporary:` directory.

When you specify the `config-file` parameter, you must specify the `commonName`, `file-name`, or `object-name` parameter.

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 `config-file` parameter specifies an OpenSSL CNF file. This file consists of sections that are delimited by a section name that is enclosed in square brackets. Each section can contain one or more properties. Some of these properties are allowed to exist outside of any section definition. As such, they are considered global. The command does not read global properties. The command reads the `[req]` section of the configuration file but actually reads only the following properties of the `[req]` section.

```
[req]
  [distinguished_name]
  [attributes]
  [req_extensions]
  [x509_extensions]
  string_mask
```

Each of these properties can define a section of its own with the exception of `string_mask` that must be either `utf8only` or `nombstr`.

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.

To delete a key, use the **no key** *keyname* command, see [“key” on page 734](#).

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.williamswillow.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 "williamswillow Australia, Ltd." OU "Customer
Support" CN www.bob.williamswillow.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 "williamswillow Australia, Ltd." OU "Customer Support"
CN www.bob.williamswillow.com.au rsa 2048 out bob
password-alias WaltzingMatilda
```

Related reference

[“password-map” on page 740](#)

This command manages the encrypted passwords to a password aliases in a password map file.

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 password-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 54 on page 742 list the available keywords. The default value is HIGH:MEDIUM:!
aNULL:!
eNULL:@STRENGTH.

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 <i>high</i> encryption cipher suites, which are currently AES or 3DES cipher suites.
MEDIUM	All <i>medium</i> encryption cipher suites, which are currently 128-bit RC2 and 128-bit RC4 cipher suites.
LOW	All <i>low</i> encryption cipher suites, which are currently 64-bit RC4 and 56-bit DES non-export cipher suites.

<i>Table 54. Available algorithm keywords for the cipher string (continued)</i>	
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 <i>man-in-the-middle</i> 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 1.0 cipher suites.
SSLv3	SSL 3.0 cipher suites.
SSLv2	SSL 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 55 on page 744 lists the available options.

<i>Table 55. SSL options as string</i>	
String value	Description
OpenSSL-default	The default value.
Disable-SSLv2 ¹	Disallows the use of SSL 2.0.
Disable-SSLv3	Disallows the use of SSL 3.0.
Disable-TLSv1	Disallows the use of TLS 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 1.1.
Disable-TLSv1d2	Disallows the use of TLS 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 3.0 and TLS 1.0, enter Disable-SSLv3+Disable-TLSv1.

clientalist {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 2.0.

```
# 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 2.0 and TLS 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
```

Related reference

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

[“idcred” on page 733](#)

This command creates identification credentials.

[“valcred” on page 749](#)

This command enters Validation Credentials mode.

[sshserverprofile](#)

This command enters SSL client profile mode.

Syntax

sshserverprofile

Guidelines

The **sshserverprofile** command enters SSH server profile mode to modify the ciphers for the SSH server profile.

Related concepts

[“SSH Server Profile commands” on page 959](#)

SSH Server Profile mode provides the commands to manage the cipher suites for SSH connections.

[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
```

Related reference

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

[“idcred” on page 733](#)

This command creates identification credentials.

ssl-client

This command enters SSL client profile mode.

Syntax

ssl-client *name*

no ssl-client *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.

Guidelines

The **ssl-client** command enters SSL Client Profile mode to create or modify an SSL client profile. An SSL client profile secures connections between the appliance and its targets.

Use the **no usergroup** command to delete an SSL client configuration.

Related concepts

[“SSL Client Profile commands” on page 959](#)

SSL Client Profile mode provides the commands to create or modify an SSL client profile.

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.

Related concepts

[“SSL Server Profile commands” on page 965](#)

You can use the SSL Server Profile commands to create or modify an SSL server profile.

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
```

Related reference

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

[“key” on page 734](#)

This command creates an alias for a private key.

[“password-map” on page 740](#)

This command manages the encrypted passwords to a password aliases in a password map file.

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.

Related reference

[“profile \(deprecated\)” on page 742](#)

This command creates a cryptographic profile that specifies an SSL service level.

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.

Related reference

[“custom” on page 751](#)

This command specifies the name of a custom time zone.

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
```

Related reference

[“daylight-start-month” on page 754](#)

This command specifies the month of the year when daylight saving time starts.

[“daylight-start-minutes” on page 753](#)

This command specifies the minute in the hour when daylight saving time starts.

[“daylight-start-week” on page 755](#)

This command specifies the instance of the day in the month when daylight saving time starts.

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
```

Related reference

[“daylight-start-day” on page 752](#)

This command specifies the day of the week when daylight saving time starts.

[“daylight-start-minutes” on page 753](#)

This command specifies the minute in the hour when daylight saving time starts.

[“daylight-start-month” on page 754](#)

This command specifies the month of the year when daylight saving time starts.

[“daylight-start-week” on page 755](#)

This command specifies the instance of the day in the month when daylight saving time starts.

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
```

Related reference

[“daylight-start-hours” on page 752](#)

This command specifies the hour in the day when daylight saving time starts.

[“daylight-start-day” on page 752](#)

This command specifies the day of the week when daylight saving time starts.

[“daylight-start-month” on page 754](#)

This command specifies the month of the year when daylight saving time starts.

[“daylight-start-week” on page 755](#)

This command specifies the instance of the day in the month when daylight saving time starts.

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
```

Related reference

[“daylight-start-day” on page 752](#)

This command specifies the day of the week when daylight saving time starts.

[“daylight-start-minutes” on page 753](#)

This command specifies the minute in the hour when daylight saving time starts.

[“daylight-start-hours” on page 752](#)

This command specifies the hour in the day when daylight saving time starts.

[“daylight-start-week” on page 755](#)

This command specifies the instance of the day in the month when daylight saving time starts.

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
```

Related reference

[“daylight-start-day” on page 752](#)

This command specifies the day of the week when daylight saving time starts.

[“daylight-start-hours” on page 752](#)

This command specifies the hour in the day when daylight saving time starts.

[“daylight-start-minutes” on page 753](#)

This command specifies the minute in the hour when daylight saving time starts.

[“daylight-start-month” on page 754](#)

This command specifies the month of the year when daylight saving time starts.

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
```

Related reference

[“daylight-stop-hours” on page 756](#)

This command specifies the hour in the day when daylight saving time ends.

[“daylight-stop-minutes” on page 757](#)

This command specifies the minute in the hour when daylight saving time ends.

[“daylight-stop-month” on page 758](#)

This command specifies the month of the year when daylight saving time ends.

[“daylight-stop-week” on page 759](#)

This command specifies the instance of the day in the month when daylight saving time ends.

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
```

Related reference

[“daylight-stop-day” on page 756](#)

This command specifies the day of the week when daylight saving time stops.

[“daylight-stop-minutes” on page 757](#)

This command specifies the minute in the hour when daylight saving time ends.

[“daylight-stop-month” on page 758](#)

This command specifies the month of the year when daylight saving time ends.

[“daylight-stop-week” on page 759](#)

This command specifies the instance of the day in the month when daylight saving time ends.

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
```

Related reference

[“daylight-stop-hours” on page 756](#)

This command specifies the hour in the day when daylight saving time ends.

[“daylight-stop-day” on page 756](#)

This command specifies the day of the week when daylight saving time stops.

[“daylight-stop-month” on page 758](#)

This command specifies the month of the year when daylight saving time ends.

[“daylight-stop-week” on page 759](#)

This command specifies the instance of the day in the month when daylight saving time ends.

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
```

Related reference

[“daylight-stop-day” on page 756](#)

This command specifies the day of the week when daylight saving time stops.

[“daylight-stop-minutes” on page 757](#)

This command specifies the minute in the hour when daylight saving time ends.

[“daylight-stop-hours” on page 756](#)

This command specifies the hour in the day when daylight saving time ends.

[“daylight-stop-week” on page 759](#)

This command specifies the instance of the day in the month when daylight saving time ends.

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
```

Related reference

[“daylight-stop-day” on page 756](#)

This command specifies the day of the week when daylight saving time stops.

[“daylight-stop-hours” on page 756](#)

This command specifies the hour in the day when daylight saving time ends.

[“daylight-stop-minutes” on page 757](#)

This command specifies the minute in the hour when daylight saving time ends.

[“daylight-stop-month” on page 758](#)

This command specifies the month of the year when daylight saving time ends.

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
```

Related reference

[“offset-hours” on page 762](#)

This command specifies the offset in hours, relative to Coordinated Universal Time, of the time zone.

[“offset-minutes” on page 762](#)

This command specifies the offset in minutes, relative to Coordinated Universal Time, of the time zone.

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.
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
```

Related reference

[“direction” on page 760](#)

This command specifies the direction, relative to Coordinated Universal Time, of the time zone.

[“offset-minutes” on page 762](#)

This command specifies the offset in minutes, relative to Coordinated Universal Time, of the time zone.

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
```

Related reference

[“direction” on page 760](#)

This command specifies the direction, relative to Coordinated Universal Time, of the time zone.

[“offset-hours” on page 762](#)

This command specifies the offset in hours, relative to Coordinated Universal Time, of the time zone.

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
```

timeout

This command sets the duration to wait for a response from an NTP server.

Syntax

timeout *milliseconds*

Parameters

milliseconds

Specifies the number of milliseconds to wait for a response from an NTP server. Enter a value in the range 750 - 3000. The default value is 750.

Guidelines

The **timeout** command sets the duration to wait for a response from an NTP server before trying the next server in the list.

Example

Identify the NTP servers, set the duration to wait for a response to 1 second, and set the clock synchronization interval to 5 minutes.

```
# ntp-service
NTP Service configuration
# remote-server Chronos-1
# remote-server Chronos-2
# timeout 1000
# 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.

force-ip-preference

This command sets whether to restrict DNS queries to the preferred IP version.

Syntax

force-ip-preference { on | off }

Parameters

on

Sends DNS queries to only the preferred IP version: A (IPv4) or AAAA (IPv6) record. This setting avoids latency with unneeded IPv6 queries in IPv4-only environments and vice versa.

off

Sends DNS queries for both A and AAAA records. This setting is the default setting.

Guidelines

The **force-ip-preference** command sets whether to restrict DNS queries to the preferred IP version. You want to force the IP preference except when both IPv4 and IPv6 addresses are in use. When not forced, the appliance resolves each host name by querying A and AAAA records and waiting for both responses or a timeout. Waiting for the response or timeout for both records can introduce unnecessary latency in DNS resolution.

To set the preferred IP version, use the **ip-preference** command.

Related reference

[“ip-preference” on page 765](#)

This command sets the preferred IP version when the DNS provider publishes both versions of addresses.

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.

Related reference

[“load-balance” on page 766](#)

This command sets the load-balancing algorithm that the appliance uses to resolve host names.

[“timeout” on page 769](#)

This command sets the time to wait before the next query 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.

Related reference

[“load-balance” on page 766](#)

This command sets the load-balancing algorithm that the appliance uses to resolve host names.

[“retries” on page 767](#)

This command sets the number of times that the appliance attempts a failed query.

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.

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 269](#) and [“Changing IP addresses in disaster recovery configurations” on page 289](#).

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
```

force-mode

This command indicates whether to force Ethernet physical mode instead of autonegotiation behavior.

Syntax

force-mode { on | off }

Parameters

on

Force the physical mode set by the **mode** command.

off

Use autonegotiation. This setting is the default value.

Guidelines

The **force-mode** command indicates whether to force Ethernet physical mode instead of autonegotiation behavior.

- By default, autonegotiation is used.
- When enabled, the physical mode is forced. No autonegotiation is done at the Ethernet driver level. Enable this option only when IBM® Support diagnoses that you are encountering a problem.

The command is not meaningful when the mode set with the mode command is Auto.

Related reference

[“mode” on page 779](#)

This command sets the interface speed and direction.

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
```

hardware-offload

This command indicates whether to offload TCP/IP packet processing.

Syntax

hardware-offload {on | off}

Parameters

on

Enables the hardware offload of TCP/IP packet processing. This setting is the default value.

off

Disables the hardware offload of TCP/IP packet processing.

Guidelines

The **hardware-offload** command indicates whether to offload TCP/IP packet processing.

- When enabled, offloads TCP/IP packet processing of Ethernet device drivers and chips. Hardware offload can improve performance.
- When disabled, does not offload TCP/IP packet processing. Disable this option only when IBM support diagnosed that you are encountering this problem.

If you disable the hardware offload and then re-enable offloading, you must restart the appliance for the change to take effect. To modify the operational behavior temporarily, use the **disable-hardware-offload** command.

Example

Disable hardware offload at the direction of IBM Support.

```
# hardware-offload off
```

Related reference

[“disable-ethernet-hardware-offload” on page 771](#)

This command manages the temporary disabling of 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
```

Related reference

[“ip-config-mode” on page 773](#)

This command identifies the configuration mode for the Ethernet interface.

[“ip-secondary-address” on page 775](#)

This command manages secondary network addresses for the Ethernet interface.

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
```

Related reference

[“ip-address” on page 772](#)

This command assigns the primary network address for the Ethernet interface.

[“ip-route” on page 774](#)

This command manages static routes in the routing table for the Ethernet interface.

[“ip-secondary-address” on page 775](#)

This command manages secondary network addresses for the Ethernet interface.

[“ipv4-default-gateway” on page 776](#)

This command designates the default IPv4 gateway for the Ethernet interface.

[“ipv6-default-gateway” on page 777](#)

This command designates the default IPv6 gateway for the Ethernet interface.

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
```

Related reference

[“ip-config-mode” on page 773](#)

This command identifies the configuration mode for the Ethernet interface.

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
```

Related reference

[“ip-address” on page 772](#)

This command assigns the primary network address for the Ethernet interface.

[“ip-config-mode” on page 773](#)

This command identifies the configuration mode for the Ethernet interface.

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.

Related reference

[“ip-config-mode” on page 773](#)

This command identifies the configuration mode for the Ethernet interface.

[“ipv6-default-gateway” on page 777](#)

This command designates the default IPv6 gateway for the Ethernet interface.

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.

Related reference

[“ipv6-nd-retransmit-timer” on page 777](#)

This command sets the interval between IPv6 neighbor discovery attempts for the Ethernet 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.

Related reference

[“ip-config-mode” on page 773](#)

This command identifies the configuration mode for the Ethernet interface.

[“ipv4-default-gateway” on page 776](#)

This command designates the default IPv4 gateway for the Ethernet interface.

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-aggregation-mode

This command manages Ethernet interfaces in the aggregate interface.

Syntax

link-aggregation-mode { on | off }

Parameters

on

Sets the interface as part of an aggregate interface.

off

Does not set the interface as part of an aggregate interface. This setting is the default value.

Guidelines

The **link-aggregation-mode** command indicates whether the interface is part of an aggregate interface. When the Ethernet interface is part of an aggregate interface, the appliance ignores configuration data about the physical Ethernet interface. In other words, you cannot modify the Ethernet interface when it is part of an aggregate interface.

Examples

Set the interface as part of an aggregate interface.

```
# link-aggregation-mode on
```

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.  
#
```

Failure notification commands

You can use the failure notification commands to control the content and destination of error reports.

The failure notification commands can be run from the command line interface in failure notification mode. To enter failure notification mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:

```
config
```

2. From global configuration mode, enter failure notification mode:

```
failure-notification
```

3. Type `exit` to leave the failure notification mode and save your changes, then type `exit` again to leave global configuration mode.

If you want to include first failure data capture information in your report, you must set **upload-report** to TRUE.

always-on-shutdown

This command indicates whether to generate an error report on shutdown.

Syntax

always-on-shutdown { on off }

on

Generate the report. This setting is the default value.

off

Does not generate the report.

Guidelines

The **always-on-shutdown** command indicates whether to generate an error report on each shutdown. Because error reports are written to the temporary: directory, you must configure failure notification to send reports as email or upload reports. The appliance deletes the contents of the temporary: directory during appliance restart.

- To send error reports as email, use the **use-smtp** command.
- To upload error reports, use the **upload-report** command.

Related reference

[“always-on-startup” on page 782](#)

This command indicates whether to send an error report on each restart.

[“upload-report” on page 786](#)

This command indicates whether to upload the error report.

[“use-smtp” on page 786](#)

This command indicates whether to use SMTP to send the error report as an email message.

always-on-startup

This command indicates whether to send an error report on each restart.

Syntax

always-on-startup { on | off }

on

Sends the report.

off

Does not automatically send the report on startup. This setting is the default value.

Guidelines

The **always-on-startup** command indicates whether to send an error report on each firmware restart. Use this command to ensure that any restart generates an error report.

email-address

This command specifies the email address to which to send error reports.

Syntax

email-address *address*

Parameters

address

Specifies the full email address of the message recipient.

Function

The **email-address** command specifies the email address to which to send error reports.

Example

Specify the email address of the recipient.

```
# email-address techsupport@example.com  
#
```

Related reference

[“email-sender-address” on page 783](#)

This command specifies the email address of the sender of the error report.

[“location-id” on page 785](#)

This command specifies the subject line of the email message.

[“remote-address” on page 785](#)

This command identifies the remote SMTP server to send the failure notification.

email-sender-address

This command specifies the email address of the sender of the error report.

Syntax

email-sender-address *address*

Parameters

address

Specifies the full email address of the message sender.

Guidelines

The **email-sender-address** command specifies the email address of the sender of the error report. If you do not use this command to specify the email address of the sender, the value set for the **email-address** is used as the sender address.

Examples

Specify the email address of the message sender.

```
# email-sender-address tech-support@example.com  
#
```

Related reference

[“email-address” on page 782](#)

This command specifies the email address to which to send error reports.

ffdc event-log

This command indicates whether to generate a background event log.

Syntax

ffdc event-log { on | off }

Parameters

on

Includes the event log.

off

Does not include the event log. This setting is the default value.

Guidelines

The **ffdc event-log** command indicates whether to run an always-on background log message capture, capturing all log and trace points within the code with minimal processor usage. Log messages that are captured are independent of the system logging configuration. This information is formatted into an error report when a first failure data capture (FFDC) event is triggered. Enable this feature to help resolve problems.

You can only use this command in conjunction with the **upload-report on** command.

ffdc memory-trace

This command indicates whether to generate a background memory leak capture.

Syntax

```
ffdc memory-trace { on | off }
```

Parameters

on

Includes the memory leak.

off

Does not include the memory leak. This setting is the default value.

Guidelines

The **ffdc memory-trace** command indicates whether to run a memory leak capture if the system memory usage is too low or too high. This option enables background leak detection, capturing all allocation call sites when the system detects a memory leak trend. This information is formatted into an error report when a first failure data capture (FFDC) event is triggered. Enable this feature to help resolve problems.

You can only use this command in conjunction with the **upload-report on** command.

ffdc packet-capture

This command indicates whether to generate a background packet capture.

Syntax

```
ffdc packet-capture { on | off }
```

Parameters

on

Includes the packet capture.

off

Does not include the packet capture. This setting is the default value.

Guidelines

The **ffdc packet-capture** command indicates whether to run an always-on background packet capture, capturing network packets on all interfaces, which include the internal loopback interface. This information is formatted into an error report when a first failure data capture (FFDC) event is triggered. Enable this feature to help resolve problems. If you enable this feature and manually request a packet capture, two packet captures are running, which significantly impacts performance.

You can only use this command in conjunction with the **upload-report on** command.

internal-state

This command indicates whether to include a snapshot of the internal state.

Syntax

```
internal-state { on | off }
```

Parameters

on

Includes the snapshot.

off

Does not include the snapshot. This setting is the default value.

Guidelines

The **internal-state** command indicates whether to include a snapshot of the internal state. Including the internal state in the error report can aid in problem determination.

location-id

This command specifies the subject line of the email message.

Syntax

location-id *string*

Parameters**string**

Specifies descriptive text.

Guidelines

The **location-id** command specifies the subject line of the email message. If the message contains spaces, wrap the value in double quotation marks.

Examples

Provide a descriptive subject line.

```
# location-id "South Campus Building 9 5th Floor"  
#
```

protocol

This command indicates the protocol to use to upload the error report.

Syntax

protocol *protocol*

Parameters**protocol**

Sets the protocol scheme to upload the error report.

v 9.1.5 mqdiag

Sends the report to the mqdiag: directory on the RAID disk.

smtp

Sends the error report to a specified email address.

temporary

Uploads the error report to a temporary local location.

remote-address

This command identifies the remote SMTP server to send the failure notification.

Syntax

remote-address *host*

Parameters

host

Identifies the remote SMTP server by host name or IP address.

Guidelines

The **remote-address** command identifies remote SMTP server to send the failure notification.

report-history

This command indicates the number of error reports to store.

Syntax

report-history *number*

Parameters

number

Identifies the number of error reports to keep. Enter a value in the range 2 - 10. The default value is 5.

Guidelines

Locally stored error reports overwrite the oldest report when the entered number is exceeded. Remotely stored error reports cannot be overwritten and are kept until manually removed.

upload-report

This command indicates whether to upload the error report.

Syntax

upload-report { on | off }

Parameters

on

Uploads the error report to a specified location.

off

Does not upload the error report. This setting is the default value.

Guidelines

You must set upload report to on if you want to include any first failure data capture (FFDC) information in your report.

use-smtp

This command indicates whether to use SMTP to send the error report as an email message.

Syntax

use-smtp { on | off }

Parameters

on

Sends the error report to a specified email address.

off

Does not send the error report. This setting is the default value.

Guidelines

You can use the **use-smtp on** command together with **upload-report on** command to enable the inclusion of first failure data capture (FFDC) information in the report.

Otherwise you can use **use-smtp on** to send a report without FFDC information when the appliance starts up or shuts down.

M 2000 M 2001 Fibre channel volume commands

You can use the fibre channel volume commands to configure fibre channel volumes on the IBM MQ Appliance.

Note: These commands are available only on M2000 and M2001 appliances, they are not available on M2002 appliances.

The fibre channel volume commands can be run from the command line interface in fibre channel volume configuration mode. To enter fibre channel volume configuration mode, complete the following steps:

1. From the appliance command line, enter global configuration mode:

```
config
```

2. From global configuration mode, enter fibre channel volume configuration mode:

```
fibre-channel-volume name
```

where *name* is the name of the volume 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.

M 2000 M 2001 `lun-uid`

This command specifies the lun-uid (LUID) used to identify the SAN storage device corresponding to the volume.

Syntax

lun-uid*lun-uid*

Parameters

lun-uid

The globally unique type 3 (FC Name_Identifier) or type 2 (IEEE EUI-64) LUID that should be used to locate the storage device. Provide the 128-bit or 64-bit value as a hexadecimal string.

Guidelines

You can discover the LUIDs that are available to the appliance by typing the following command:

```
show fibre-channel-luns
```

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

M2000 M2001 *use-multipath*

This command indicates whether multipath connections are enabled for a volume connecting to a SAN.

Syntax

use-multipath {on | off}

Parameters

on

Enables multipath connections. This setting is the default value.

off

Disables multipath connections.

Guidelines

Typically, SAN infrastructures have multiple routes configured in the SAN fabric to each SAN host. Using multipath combines these routes to appear as a single route to the queue manager, but with the underlying advantages of reliability and performance offered by multiple routes. The multipath option is selected by default for a volume. If you know that you do not have multiple physical routes configured, select off.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

Example

Disable multipath connections.

```
# use-multipath off
```

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.wsd1 local:///stock.wsd1 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.xml
store:///InitialConvert.xml
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/InitialConvert.xml
store:///InitialConvert.xml
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.

acl

This command enters Access Control List mode.

Syntax

Creates or edits a service-specific ACL.

```
acl name
```

Edits the ACL for the SSH service.

```
acl ssh
```

Edits the ACL for the web management interface.

```
acl web-mgmt
```

Edits the ACL for the XML management interface.

```
acl xml-mgmt
```

Deletes a service-specific ACL.

```
no acl 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.

ssh

Identifies the SSH service. On XI50z, the dp-admin account cannot edit the ACL.

web-mgmt

Identifies the web management interface. On XI50z, the dp-admin account cannot edit the ACL.

xml-mgmt

Identifies the XML management interface. On XI50z, the dp-admin account cannot edit the ACL.

Guidelines

The **acl** command enters Access Control List mode. In this mode, you can configure an ACL for a specific service.

An ACL contains one or more clauses. Each clause consists of an IP address range that is defined by an IP address and netmask and a Boolean value (allow or deny). IP addresses are evaluated against each clause in the order in which they are in the list. A candidate address is denied or granted access to the service with the first matching clause. Therefore, the order of clauses is important.

Use the **no acl** command to delete a service-specific ACL.

Examples

- Create the ACL -1 ACL.

```
# acl ACL-1
```

- Delete the ACL -1 ACL.

```
# no acl ACL-1
```

- Edit the ACL for the SSH service.

```
# acl ssh
```

- Edit the ACL for the web management interface.

```
# acl web-mgmt
```

- Edit the ACL for the XML management interface.

```
# acl xml-mgmt
```

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.

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
#
```

crypto

This command enters Crypto mode.

Syntax

crypto

Guidelines

The **crypto** command enters Crypto mode.

Use the **exit** command to exit Crypto mode and return to Global mode.

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
```

Related concepts

[“DNS commands” on page 764](#)

You can use the DNS commands to configure the DNS settings on the IBM MQ Appliance.

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-re-create 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 disabled
# exit
```

- Delete the configuration for the eth5 Ethernet interface.

```
# no ethernet eth5
```

Related concepts

[“Ethernet commands” on page 769](#)

You can use the Ethernet commands to configure the Ethernet interfaces on the IBM MQ Appliance.

failure-notification

This command enters Failure Notification mode.

Syntax

failure-notification

no failure-notification

Guidelines

Use the **no failure-notification** command to disable failure reporting. By default, failure reporting is disabled.

Related concepts

[“Failure notification commands” on page 781](#)

You can use the failure notification commands to control the content and destination of error reports.

flash

This command enters Flash mode.

Syntax

flash

Guidelines

The **flash** command enters Flash mode. In Flash mode, use the available command to manage firmware images and the files on the appliance.

To exit Flash mode and return to Global mode, use the **exit** command.

Related concepts

[“Initialization commands” on page 823](#)

You can use the initialization commands to initialize or reset the IBM MQ Appliance.

▶ M2000 M2001 ***fibre-channel-fs-init***

This command initializes the file space for the specified volume.

Syntax

fibre-channel-fs-init *volume_name*

Parameters

volume_name

The volume to initialize the file space for.

Guidelines

You must only initialize the volume file system once, after you create it and before you attempt to use it. Initializing after you have used the volume erases all the contents of the volume.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

▶ M2000 M2001 ***fibre-channel-fs-repair***

This command repairs the file space for the specified volume.

Syntax

fibre-channel-fs-repair *volume_name*

Parameters

volume_name

The volume to repair the file space for.

Guidelines

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

▶ M2000 M2001 ***fibre-channel-unlock-volume***

This command unlocks the specified volume.

Syntax

fibre-channel-unlock-volume *volume_name*

Parameters

volume_name

The volume to unlock.

Guidance

You might require to unlock a volume, for example, if it has been left locked by an appliance that stopped abruptly while having the volume enabled. You can unlock the volume from another appliance to take over the work from the failed appliance.

The appliance that unlocks a volume must be zoned such that it can see the volume.



Attention: You should not use this command other than in fault situations. If you unlock a volume currently in use by another appliance with no fault, you will cause the queue manager to crash.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

M 2000 M 2001 *fibres-channel-volume*

This command creates a volume used to connect to a SAN using the appliance fibre channel.

Syntax

Creates or modifies a volume.

fibres-channel-volume *volume*

Deletes a volume.

no fibres-channel-volume *volume*

Availability

Physical appliances only.

Parameters

volume

Specifies the name of the volume.

Guidelines

The **fibres-channel-volume** command enters volume configuration mode. While in the mode, define a SAN volume.

Use the **no fibres-channel-volume** command to remove a volume.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

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.

Related concepts

[“Host Alias commands” on page 823](#)

Host Alias mode provides the commands to create or modify a host alias.

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.

Related concepts

[“IPMI LAN channel commands” on page 831](#)

You can use the IPMI LAN channel commands to create or modify an IPMI LAN channel on the IBM MQ Appliance.

Related reference

[“ipmi-user” on page 800](#)

This command enters IPMI User mode.

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.

Related concepts

[“IPMI user commands” on page 829](#)

You can use the IPMI user commands to create or modify an IPMI users on the IBM MQ Appliance.

Related reference

[“ipmi-lan-channel” on page 799](#)

This command enters IPMI LAN Channel mode.

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
p80Jl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfR1
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi01ud5P17+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
p80Jl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfR1
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi01ud5P17+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
p80Jl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfR1
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi01ud5P17+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
```

ldap-search-parameters

This command enters LDAP Search Parameters mode.

Syntax

ldap-search-parameters *name*

no ldap-search-parameters *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.

Guidelines

The **ldap-search-parameters** command enters LDAP Search Parameters mode. In this mode, you can create or modify LDAP search parameters. These parameters are used to search an LDAP search to retrieve the distinguished name (DN) of the user.

Use the **no ldap-search-parameters** command to delete LDAP search parameters.

Related concepts

[“LDAP Search Parameters commands” on page 835](#)

LDAP Search Parameters mode provides the commands to modify LDAP search parameters.

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.

If you want to use the aggregate interface with a floating IP address for a high availability configuration, then your name should observe the following restrictions:

- The name must be 12 characters or less.
- The name 'lo' is not permitted
- The name cannot start with any of the following strings:
 - eth
 - mgt
 - vlan

For example, the names `halink`, `floatingIP`, and `S23HA` are all valid names. The names `other_ha_link`, `ethfloating`, and `mgtlink` are invalid. You can use a name that meets these restrictions with the **sethaint** command as the local interface parameter, to specify that the aggregate interface can be used to connect to the queue manager on the two appliances in the high availability group (see [“sethaint” on page 677](#)).

If you do not require to use the aggregated interface with a floating IP, then the name rules are more relaxed:

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.

Related concepts

[“Link aggregation commands” on page 837](#)

You can use the link aggregation commands to configure link aggregation interfaces on the IBM MQ Appliance.

loadbalancer-group

This command enters Load Balancer Group mode.

Syntax

loadbalancer-group *name*

no loadbalancer-group *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.

Guidelines

The **loadbalancer-group** command enters Load Balancer Group mode. After you complete the configuration of a load balancer group, assign it to an XML manager. The assignment of the load balancer group to an XML manager makes the group available to DataPower services that this XML manager supports.

Use the **no loadbalancer-group** command to delete a load balancer group.

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.

Related concepts

[“Log target commands” on page 855](#)

Log target mode provides the commands to create or modify a log target.

logsize

This command sets the size of a basic event log.

Syntax

logsize *KB*

Parameters

KB

Specifies the size in KB of the log. The default value is 200.

Guidelines

Without an argument, the **logsize** command displays the size of the log file in KB.

You must be logged in as the admin user to use this command.

Note:

Use the **loglevel**, **logsize**, and **syslog** commands to configure the basic logging system.

Use the **logging target** command to enter Logging mode. From this mode, define more precise control over log formats and contents.

Examples

- Set the log size to 250 KB.

```
# logsize 250
```

- Display the configured log size in KB.

```
# logsize
```

mkdir

This command creates a subdirectory.

Syntax

```
mkdir local:///subdirectory
```

Parameters

local:///subdirectory

The subdirectory to create in the `local:` directory.

Guidelines

The **mkdir** command creates subdirectories in the `local:` directory on the DataPower appliance. You can create subdirectories for application-specific files such as style sheets and schemas.

Use the **rmdir** command to delete subdirectories.

Examples

- Create the `stylesheets` subdirectory of the `local:` directory.

```
# mkdir local:///stylesheets
Directory 'local:///stylesheets' successfully created.
#
```

- Create the `C-1` subdirectory in the `stylesheets` subdirectory of the `local:` directory.

```
# mkdir local:///stylesheets/C-1
Directory 'local:///stylesheets/C-1' successfully created.
#
```

Related reference

[“rmdir” on page 814](#)

This command removes a subdirectory.

network

This command enters Network Settings mode or resets network settings.

Syntax

Enters Network Settings mode

```
network
```

Resets network settings

```
no network
```

Guidelines

The **network** command enters Network Settings mode. While in this mode, you can define network settings that include the following functions:

- Control the generation of certain Internet Control Message Protocol (ICMP) replies and how to manage failures. By default, the appliance replies to the corresponding ICMP requests.

- Control routing behavior.
- Control interface isolation.
- Define ECN settings.

The **no network** command resets network settings to their default settings.

Related concepts

[“Network commands” on page 911](#)

You can use the network commands to modify network settings on the IBM MQ Appliance.

ntp-service

This command enters NTP service mode.

Syntax

ntp-service

no ntp-service

Guidelines

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.

Related concepts

[“Date, time, and locale configuration commands” on page 750](#)

You can use these commands to configure the time, date, and locale on the IBM MQ Appliance.

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-initialize

This command initializes an array volume.

Syntax

raid-initialize *name*

Parameters

name

Specifies the name of the array volume. The name is `raid0`.

Guidelines

The **raid-initialize** command makes the two disks into an array volume. This action destroys any prior content of the array volume.

Example

Build a RAID volume on the disks on the system.

```
# raid-initialize 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

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.

rbm

This command enters RBM Settings mode.

Syntax

rbm

no rbm

Guidelines

While in RBM mode, you configure role-based management (RBM) settings.

Use the **no rbm** command to disable RBM service. Note that this command can disable GUI access.

Related concepts

[“RBM setting commands” on page 917](#)

After you delete the known host from the file or registry, the next attempt to connect to the appliance displays a prompt similar to the following. When prompted, enter yes to continue.

```
The authenticity of host 'Midtownjojo (xxx.xxx.xxx.xxx)' cannot be established.  
RSA key fingerprint is xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx.  
Are you sure you want to continue connecting (yes/no)? yes  
Warning: Permanently added 'Midtownjojo,xxx.xxx.xxx.xxx (RSA)' to the list of known hosts.  
(unknown)  
Unauthorized access prohibited.  
login:
```

Delete and re-create new SSH host keys for the SSH service on the appliance. When you run the command, you are presented with a prompt where you must enter y to continue.

```
# reset-ssh-keys  
Resetting SSH keys will regenerate host keys but it may terminate active SSH sessions  
and will require authenticity verification next time a client connects over SSH.  
Do you want to continue? Yes/No [y/n]: y
```

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 reenables 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.
```

rest-mgmt

This command enters REST Management Interface mode to manage the configuration of the REST management interface.

Guidelines

The **rest-mgmt** command enters REST Management Interface mode to configure the REST management interface. The REST management interface monitors REST management traffic. When enabled, You can send requests to the REST management interface to supported service protocols to manage the

appliance. The REST management interface runs SSL and uses HTTP Basic Authentication (user name and password).

Use the **no rest-mgmt** command to disable the REST management interface.

Examples

Enter REST Management Interface mode to configure the REST management interface.

```
# rest-mgmt
Modify REST Management Interface configuration
#
```

Disable the REST object management interface.

```
# no rest-mgmt
REST management: successfully disabled
#
```

Related concepts

[“REST Management Interface commands” on page 950](#)

Use the REST Management Interface commands to enable and modify the configuration of the REST management interface.

rmdir

This command removes a subdirectory.

Syntax

```
rmdir local:///subdirectory
```

Parameters

local:///subdirectory

The subdirectory to remove from the local: directory.

Guidelines

The **rmdir** command removes subdirectories from the local: directory.

Example

Deletes the stylesheets subdirectory and all its contents from the local: directory.

```
# rmdir local:///stylesheets
Removing 'local:///stylesheets' will delete all files including subdirectories!
Do you want to continue? [y/n]:y
Directory 'local:///stylesheets' successfully deleted.
#
```

Related reference

[“mkdir” on page 806](#)

This command creates a subdirectory.

save error-report

This command creates an error report.

Syntax

```
save error-report
```

Guidelines

The **save error-report** command creates an error report in the temporary:///error-report.txt file. This file can then be used for debugging purposes. The file can be copied off the appliance with the **send error-report** command.

This command takes no arguments. It automatically creates a file that contains any backtrace or watchdog error report, the contents of the audit log, and the running configuration.

If there is insufficient space to write the file, the following error message indicates this condition:

```
"Could not write error report to file."v
```

If you receive this message, delete files to make sufficient space.

Examples

Generate an error report file in temporary:///error-report.txt. This file is then sent to supportteam@example.com on the specified mail server.

```
# save error-report
# send error-report smtp.example.com "Important information" supportteam@example.com
#
```

Related reference

[“send error-report” on page 815](#)

This command sends an error report as email.

send error-report

This command sends an error report as email.

Syntax

```
send error-report mail-server subject email-address email-sender-address
```

Parameters

mail-server

Identifies a local SMTP server by IP address or by host name.

subject

Specifies the text string for the subject field of the outgoing message.

email-address

Specifies the full email address of the message recipient.

email-sender-address

Specifies the full email address of the message sender.

Guidelines

With the **send error-report** command, keep in mind the following points.

- With automated notification enabled, the appliance searches for a backtrace file (that contains failure-related data) at each restart. With a backtrace file, the appliance sends the file to the appliance but does not send any notifications.
- Place this command within the configuration file to ensure that it is run on each restart.
- To identify the SMTP server by host name, you must place this command after the **ip name-server** command in the configuration file. Otherwise, the appliance cannot resolve the host name to IP address.
- This command can be run to immediately send an error report that was generated with the **save error-report** command.

Example

On system restart, send a notification to supportteam@customer.com and weekendcoverage@attbb.com through the specified mail server.

```
# send error-report
smtp.customer.com ALERT! supportteam@customer.com weekendceverage@attbb.com
```

Related reference

[“failure-notification” on page 796](#)

This command enters Failure Notification mode.

[“save error-report” on page 814](#)

This command creates an error report.

snmp

This command enters SNMP Settings mode.

Syntax

snmp

no snmp

Guidelines

While in SNMP Settings mode, you configure Simple Network Management Protocol (SNMP) settings.

Use the **no snmp** command to disable SNMP.

Examples

- Enter SNMP Settings mode to manage the SNMP log target.

```
# snmp
SNMP Settings mode
#
```

- Disable SNMP.

```
# no snmp
#
```

Related concepts

[“SNMP Settings commands” on page 952](#)

SNMP Settings mode provides the commands to modify the SNMP settings.

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
```

system

This command enters System Settings mode.

Syntax

system

timezone

This command enters Timezone mode.

Syntax

timezone

Guidelines

While in Timezone mode, configure the time zone settings for the appliance. The time zone alters the display of time to the user.

Related concepts

[“Date, time, and locale configuration commands” on page 750](#)

You can use these commands to configure the time, date, and locale on the IBM MQ Appliance.

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
#
```

undo

This command reverts a modified configuration to its previously saved configuration.

Syntax

undo *type name*

Parameters

type

Specifies the type of configuration. For a complete list, use the **show** command.

name

Specifies the name of the configuration.

Guidelines

The **undo** command reverts a modified configuration to its last persisted state. The persisted state is the configuration in the startup configuration. You use the **write memory** command to save configuration changes to the startup configuration.

For a modified configuration in the startup configuration, you receive the following message.

```
type name - Configuration reverted.
```

For a new configuration that in not saved to the startup configuration, you receive the following message.

```
Cannot undo new configuration
```

For a configuration that is not modified, you receive the following message.

```
Cannot undo - configuration was not modified
```

For a nonexistent configuration, you receive the following message.

```
Cannot undo last configuration change
```

For a nonexistent type, you receive the following message.

```
Invalid class
```

user

This command enters User mode.

Syntax

user *name*

no user *name*

Parameters

name

Specifies the name of the user.

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 **user** command is available in Global mode. The **user** command enters User mode. While in User mode, you can create or modify User objects.

Use the **no user** command to delete a user account.

To confirm the account, use the **show usernames** command.

Related concepts

[“Appliance user commands” on page 719](#)

You can use the appliance user commands to configure appliance users on the IBM MQ Appliance.

[“Appliance user group commands” on page 723](#)

You can use the appliance user group commands to configure appliance user groups on the IBM MQ Appliance.

usergroup

This command enters User Group mode.

Syntax

usergroup *name*

no usergroup *name*

Parameters

name

Specifies the name of the user group.

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

A user group consists of a set of access privileges that are subsequently assigned to individual user accounts.

Use the **no usergroup** command to delete a user group.

Related concepts

[“Appliance user group commands” on page 723](#)

You can use the appliance user group commands to configure appliance user groups on the IBM MQ Appliance.

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
```

vlan

This command enters VLAN mode to manage the configuration of VLAN interfaces. Before you configure an interface, obtain the essential network data from your network team.

Syntax

Enters the mode to create or modifies a VLAN interface configuration

vlan *name*

Deletes a VLAN interface configuration

no vlan *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.

Guidelines

The **vlan** command enters VLAN mode to manage the configuration of VLAN interfaces. You can create multiple VLAN interface configurations on a single Ethernet interface or aggregate interface.

VLAN packets are identified by the IEEE 802.1Q protocol.

Use the **no vlan** command to delete the configuration for a VLAN interface.

Related concepts

[“VLAN Commands” on page 987](#)

You can use the VLAN commands to configure the VLAN interfaces on the IBM MQ Appliance.

web-mgmt

This command manages access to the web management service.

Syntax

Enters Web Management Service mode.

```
web-mgmt
```

Starts the web management service on a listening address-port pair with an optional session timer.

```
web-mgmt address port [ timer | off ]
```

Disables the web management service.

```
no web-mgmt
```

Parameters

address

Identifies the listening IP address on the appliance. The default value is 0.0.0.0.

- 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: For a management service, the value of 0.0.0.0 or : : is a security risk.

port

Identifies the listening port on the appliance. The default value is 9090.

timer

Sets the idle-session timer in seconds. Enter a value in the range 0 - 65535. The default value is 600. A value of 0 disables the session timer.

off

Restores the idle-session timer to its default value.

Guidelines

The **web-mgmt** command manages access to the web management service.

Use the **no web-mgmt** command to disable the web management service.

Examples

- Enters Web Management Service mode.

```
# web-mgmt
Modify Web Management Service configuration
```

- Modify the web management service for the specified IP address-port pair. The idle-session timer uses the default value of 10 minutes.

```
# web-mgmt 10.10.13.31 9090
Web management successfully started
```

- Modify the web management service for the specified IP address-port pair. Set the idle-session timer to 15 minutes.

```
# web-mgmt 10.10.13.31 9090 900
Web management successfully started
```

- Restore the idle-session timer to its default value.

```
# web-mgmt 10.10.13.31 9090 off
Web management successfully started
```

- Disable the idle-session timer.

```
# web-mgmt 10.10.13.31 9090 0
...
Web management successfully started
```

- Disable the web management service.

```
# no web-mgmt
Web management successfully disabled
```

- Enable the web management service with its previously defined settings.

```
# web-mgmt
Modify Web Management Service configuration
# admin-state enabled
# exit
```

Related concepts

[“Web management service commands” on page 997](#)

You can use the web management service commands to configure the web management service settings on the IBM MQ Appliance.

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.

Command	Purpose
“admin-state” on page 708	Sets the administrative state for the configuration.
“ip-address” on page 823	Creates an alias for an IP address on a Ethernet port.

Related reference

[“host-alias” on page 798](#)

This command enters Host Alias mode to map an IP address to an alias.

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 initialize 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
#
```

Related reference

[“boot image” on page 825](#)

This command designates the startup firmware image and restarts the appliance with this image.

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
```

Related reference

[“boot image” on page 825](#)

This command designates the startup firmware image and restarts the appliance with this image.

[“boot switch” on page 826](#)

This command switches between primary and secondary installations.

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.
```

Related reference

[“boot config” on page 824](#)

This command designates the startup configuration for the next restart.

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
```

Related reference

[“boot delete” on page 824](#)

This command deletes the secondary installation.

[“boot image” on page 825](#)

This command designates the startup firmware image and restarts the appliance with this image.

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.
```

Related reference

[“boot config” on page 824](#)

This command designates the startup configuration for the next restart.

[“shutdown” on page 828](#)

This command restarts or shuts down the appliance.

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.

You must be logged in as the user `admin` to use the **reinitialize** command.

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.scrpyt2
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 `halt` [*seconds*] (deprecated)

shutdown `poweroff` [*seconds*]

Parameters

reboot

Shuts down and restarts the appliance.

halt

Shuts down the appliance without restarting. The power to the appliance remains on. This keyword is deprecated. Use **poweroff** instead.

poweroff

Stops the appliance and turns off the power.

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 after waiting ten seconds.

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 and turn off the appliance.

```
# shutdown poweroff 60
Shutdown in 60 second(s).
#
```

Related reference

[“boot config” on page 824](#)

This command designates the startup configuration for the next restart.

verify-firmware

This command runs the integrity verifier tool against the current firmware image.

Syntax

verify-firmware

Guidelines

The **verify-firmware** command runs the integrity verifier tool against the current firmware image. This tool verifies the integrity of files that are part of the currently installed firmware.

When you run the integrity verifier tool, it verifies the integrity of files that are part of the currently installed firmware. When files are different than the currently installed firmware image, messages are written to the console and also written as warnings to the system log.

This command is similar the boot image command with the file-integrity checker.

Example

Runs the integrity verifier tool.

```
# verify-firmware
Verifying manifest signature
Examining files...
Examined 1000 records, checked 955 files
File store:AAAInfo.xml has changed
...
Examined 6000 records, checked 4807 files
Verification complete. Examined 6999 files
Checked 5145 files, 1 changed
Firmware verification completed with warnings
```

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.

Related reference

[“user-id” on page 830](#)

This command sets the identifier for the IPMI user.

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.

Related reference

[“password” on page 830](#)

This command sets the password for the IPMI user.

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.



CAUTION: Enabling IPMI over LAN exposes security vulnerability [CVE-2013-4786](#), therefore IPMI over LAN must only be enabled in a secure environment.

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.

Related reference

[“ipmi-user” on page 800](#)

This command enters IPMI User mode.

[“maximum-channel-privilege-level” on page 833](#)

This command sets the maximum privilege level for users.

[“sol-required-user-privilege-level” on page 834](#)

This command sets the privilege level for serial over LAN.

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.

Related reference

[“ip address” on page 832](#)

This command sets IP addresses with subnet mask for the IPMI LAN channel.

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.

Related reference

[“sol-required-user-privilege-level” on page 834](#)

This command sets the privilege level for serial over LAN.

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.

Related reference

“sol-enabled” on page 834

This command indicates whether to support serial over LAN.

LDAP Search Parameters commands

LDAP Search Parameters mode provides the commands to modify LDAP search parameters.

To enter the mode, use the Global **ldap-search-parameters** command.

While in this mode, use the following commands in the following table to define the LDAP search parameters.

- 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.

base-dn

This command specifies the base DN to begin the search.

Syntax

base-dn *DN*

Parameters***DN***

Specifies the base DN for the search.

Guidelines

The **base-dn** command specifies the distinguished name (DN) relative to the LDAP search. This value identifies the entry level of the tree that is used by the **scope** command.

Related reference

“scope” on page 837

This command indicates the depth of the LDAP search.

filter-prefix

This command specifies the prefix of the LDAP filter expression.

Syntax

filter-prefix *prefix*

Parameters

prefix

Specifies the prefix of the filter expression.

Guidelines

The **filter-prefix** command specifies the string prefix to construct an LDAP filter expression, as defined in RFC 4515. This string is added before the user name to construct the LDAP filter to search for the DN of the user.

If the prefix is (&(mail= and the user name is bob@example.com and the suffix is)(c=US)), the LDAP search filter is (&(mail=bob@example.com)(c=US)).

You can use the **filter-suffix** to append a string to the LDAP filter expression to complete the search filter.

Example

Create the LDAP filter expression (&(mail=bob@example.com)(c=US)) based on bob@example.com as the user name.

```
# filter-prefix "&(mail="
# filter-suffix ")(c=US))"
```

Related reference

[“filter-suffix” on page 836](#)

This command specifies the suffix of the LDAP filter expression.

Related information

[RFC 4515, Lightweight Directory Access Protocol \(LDAP\): String Representation of Search Filters](#)

filter-suffix

This command specifies the suffix of the LDAP filter expression.

Syntax

filter-prefix *suffix*

Parameters

suffix

Specifies the suffix of the filter expression.

Guidelines

The **filter-suffix** command specifies the string suffix to construct an LDAP filter expression, as defined in RFC 4515. This string is added after the user name to construct the LDAP filter to search for the DN of the user.

If the prefix is (&(mail= and the user name is bob@example.com and the suffix is)(c=US)), the LDAP search filter is (&(mail=bob@example.com)(c=US)).

You must use the **filter-prefix** to add the prefix string to the LDAP filter expression to complete the search filter.

Example

Create the LDAP filter expression (&(mail=bob@example.com)(c=US)) based on bob@example.com as the user name.

```
# filter-prefix "&(mail="
# filter-suffix ")(c=US))"
```

Related reference

[“filter-prefix” on page 835](#)

This command specifies the prefix of the LDAP filter expression.

Related information

[RFC 4515, Lightweight Directory Access Protocol \(LDAP\): String Representation of Search Filters](#)

returned-attribute

This command specifies the attribute to return for each match.

Syntax

returned-attribute *attribute*

Parameters

attribute

Specifies the name of the attribute to return. The default value is dn.

Guidelines

The **returned-attribute** command specifies the name of the attribute to return for each entry that matches the search criteria.

scope

This command indicates the depth of the LDAP search.

Syntax

scope { base | one-level | subtree }

Parameters

base

Searches the entry level of the tree only.

one-level

Searches the entry level of the tree and any object that is one-level below the input.

subtree

Search the entry level of the tree and all of its descendants. This setting is the default value.

Guidelines

The **scope** command indicates the depth of the LDAP search. The entry level of the tree is defined by the **base-dn** command.

Related reference

[“base-dn” on page 835](#)

This command specifies the base DN to begin the search.

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.

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
```

Related reference

[“ip-config-mode” on page 839](#)

This command identifies the configuration mode for the aggregate interface.

[“ip-secondary-address” on page 841](#)

This command manages secondary network addresses for the aggregate interface.

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
```

Related reference

[“ip-address” on page 838](#)

This command assigns the primary network address for the aggregate interface.

[“ip-route” on page 840](#)

This command manages static routes in the routing table for the aggregate interface.

[“ip-secondary-address” on page 841](#)

This command manages secondary network addresses for the aggregate interface.

[“ipv4-default-gateway” on page 842](#)

This command designates the default IPv4 gateway for the aggregate interface.

[“ipv6-default-gateway” on page 843](#)

This command designates the default IPv6 gateway for the aggregate interface.

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
```

Related reference

[“ip-config-mode” on page 839](#)

This command identifies the configuration mode for the aggregate interface.

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
```

Related reference

[“ip-config-mode” on page 839](#)

This command identifies the configuration mode for the aggregate interface.

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.

Related reference

[“ip-config-mode” on page 839](#)

This command identifies the configuration mode for the aggregate interface.

[“ipv6-default-gateway” on page 843](#)

This command designates the default IPv6 gateway for the aggregate interface.

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.

Related reference

[“ipv6-nd-retransmit-timer” on page 843](#)

This command sets the interval between IPv6 neighbor discovery attempts for the aggregate interface.

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.

Related reference

[“ip-config-mode” on page 839](#)

This command identifies the configuration mode for the aggregate interface.

[“ipv4-default-gateway” on page 842](#)

This command designates the default IPv4 gateway for the aggregate interface.

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 bandwidth
# lACP-hash mac
```

Related reference

[“type” on page 847](#)

This command defines the mode to use for link aggregation.

lACP-select

This command sets the algorithm for the LACP selection policy.

Syntax

```
lACP-select { bandwidth | count | stable }
```

Parameters

bandwidth

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 bandwidth
# lacp-hash mac
```

Related reference

[“type” on page 847](#)

This command defines the mode to use for link aggregation.

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
```

MTU

This command sets the maximum transmission unit of the aggregate 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 aggregate interface. The MTU of the aggregate interface overrides the MTU of the Ethernet interface.

The MTU is determined regardless of the length of the layer 2 encapsulation.

Example

Set the MTU to 4 KB.

```
# mtu 4096
```

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 bandwidth
# lacp-hash mac
```

Related reference

[“lacp-hash” on page 843](#)

This command sets which hash function the LACP aggregation uses to determine the interface for outbound packets.

[“lacp-select” on page 844](#)

This command sets the algorithm for the LACP selection policy.

yield-standby

This command manages the inclusion of the aggregate interface in the standby group.

Syntax

yield-standby

Guidelines

The **yield-standby** command removes this appliance in the standby group. This command does not modify the interface configuration. After removal, the interface is added back to the standby group in potentially a different state.

Use this command after you quiesce the appliance to remove the interface gracefully from its standby group. You quiesce the appliance for maintenance.



Attention: Never use this command when preemptions is enabled in the standby control configuration.

The state of the interface in the standby group controls what happens after you run the command.

- When you issue the **yield-standby** command against the active appliance, the following changes occur.
 - The active appliance resigns from the standby group and a standby control takeover occurs, which potentially breaks in flight connections and transactions.
 - The standby appliance becomes the active appliance.
 - An election occurs among the group members to determine which one becomes the standby appliance.
- When you issue the **yield-standby** command against the standby appliance, the following changes occur.
 - The standby appliance is temporarily removed from the standby group.
 - An election occurs among the group members to determine which one becomes the standby appliance.
- When you issue the **yield-standby** command against a listen appliance, that appliance is removed and added back without an observable difference.

Example

Temporarily remove the appliance from the standby group.

```
# yield-standby
```

Load Balancer Group commands

Load Balancer Group mode provides the commands to create or modify a load balancer group configuration.

To enter the mode, use the Global **loadbalancer-group** command. To delete a load balancer group, use the Global **no loadbalancer-group** command.

While in this mode, use the following commands to define a load balancer group.

- 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.

algorithm

This command specifies the server selection algorithm.

Syntax

```
algorithm{ first-alive | hash | least-connections | round-robin | weighted-least-connections | weighted-round-robin }
```

Parameters

first-alive

Uses the concept of a primary server and 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 backup servers. The primary server is the first server in the members list.

hash

Uses the IP address of the client as the basis for server selection.

With the hash algorithm, the same client is served by the same server. Use this algorithm only when clients access applications that require the storage of server-side state information, such as cookies. Hashing algorithms cannot ensure even distribution.

least-connections

Maintains a record of active server connections and forward a new connection to the server with the least number of active connections.

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.

weighted-round-robin

Maintains a weighted list of servers and forwards new connections in proportion to the weight (or preference) of each server.

Guidelines

The **algorithm** command specifies the server selection algorithm. A request to connect to a load balancer group results in a server with a health state of up selected from the pool according to the algorithm. The algorithm provides a method for selecting which server with a health state of up receives an incoming client request.

Examples

- Specify that server selection uses the *first-alive* algorithm.

```
# algorithm first-alive
```

- Specify that server selection uses the least-connections algorithm.

```
# algorithm least-connections
```

- Specify that server selection uses the weighted-round-robin algorithm.

```
# algorithm weighted-round-robin
```

damp

This command specifies the dampening period for a server with health state of `softdown`.

Syntax

damp *seconds*

Parameters

seconds

Specifies the number of seconds that a server remains in a `softdown` state. Enter a value in the range 1 - 86400. The default value is 120.

Guidelines

The **damp** command specifies the dampening period for a member server. The dampening period is the duration that a server is removed from the load balancer group because it cannot connect during a normal HTTP or TCP transaction. Such a server has a health state of `softdown`. When this interval expires, the server is restored to the load balancer group and placed in the up state.

This command does not affect servers that are in the down state.

Example

Set the dampening period of 5 minutes.

```
# damp 600  
#
```

Related reference

[“algorithm” on page 849](#)

This command specifies the server selection algorithm.

[“server” on page 854](#)

This command adds a member to the load balancer group.

giveup-when-all-members-down

This command specifies the connection-behavior when no member is up.

Syntax

giveup-when-all-members-down { on | off }

Parameters

on

Does not forward the connection to any member. Makes the next attempt when at least one member is in the up state.

off

Selects the first member in the down state and forwards the connection to this server. This setting is the default value.

Examples

- Disable connection attempts if all members are in the down state.

```
# giveup-when-all-members-down on
#
```

- Restore the default state.

```
# giveup-when-all-members-down off
#
```

health-check

This command defines the periodic health check procedure.

Syntax

health-check *admin-state target-uri target-port type use-soap send-soap timeout frequency xpath filter [sslproxy] [enforce-timeout] [independent-checks] [gatewayscript-checks] [request-method] [request-custom-method] [request-doc] [request-content-type] [response-evaluator-metadata] [response-evaluator]*

Parameters

admin-state

Controls whether to run a periodic health check.

on

Enables the health check.

off

Disables the health check. This setting is the default state.

target-uri

For a standard health check, specifies the non-server (file path) portion of the target URI. That is, specify the URI to receive the client request that the rule generates. The default value is /.

This URI is used with the specified remote port.

target-port

Specify the port on the target server to receive the query. The default value is 80.

You can override this value for one or more members of the Load Balancer Group with the *health-port* argument of the **server** command.

The response from the server is evaluated to determine the health status of each member server in the group. The request is sent to the target URI and remote port.

This port is used for LDAP and standard health checks.

type

Controls the type of check.

Standard

Checks the health with an HTTP request on the remote port. The port is specified by the *port* argument unless it is overridden by members of the load balancer group with the **health-port** argument of the **server** command. The standard setting is the default value.

TCPConnection

Checks the health with a TCP connection request on the remote port. The port is specified by the *port* argument unless it is overridden by members of the load balancer group with the *health-port* argument of the **server** command.

use-soap

For a standard health check, specifies the HTTP method to access the target URI.

on

Accesses the target URI with an HTTP POST operation by posting a SOAP message. This setting is the default value.

off

Accesses the target URI with an HTTP GET operation.

send-soap

When the *use-SOAP* argument is on, specify the SOAP message to send as a client request. The default value is `store:///healthcheck.xml`. When the *use-SOAP* argument is off, use two double quotation marks.

timeout

Specifies the number of seconds for the completion of the health check. Enter a value in the range 2 - 86400. The default value is 10.

If successful, the server is deemed healthy and is marked as up; otherwise, the server is marked as down.

frequency

Specifies the number of seconds between health checks. Enter a value in the range 5 - 86400. The default value is 180.

xpath

Use with the *filter* argument to specify the XPath expression that must be found in a valid server response.

filter

Specifies the style sheet to filter the server response. The default value is `store:///healthcheck.xsl`.

This style sheet uses the specified *xpath* argument as input and scans the server response for its presence. If found, the server is deemed healthy and is marked as up; otherwise, the server is marked as down.

sslproxy

For a standard health check, specifies the name of the SSL Proxy Profile to secure the connection.

enforce-timeout

For a standard health check, specifies whether to use the health check timeout value to interrupt and end a health check transaction.

on

Specifies that the health check timeout value is used.

off

Specifies that the health check timeout value is not used. This setting is the default value.

independent-checks

For a standard health check, specifies whether the health check transactions in a Load Balancer Group run independently or sequentially.

on

Specifies that the health check transactions run independently.

off

Specifies that the health check transactions run sequentially. This setting is the default value.

Guidelines

A health check is a scheduled rule that sends the same request to each member. The successful completion of the health check requires that the server passes normal TCP and HTTP connection criteria, depending on check type. Optionally, a standard health check can use a filter to evaluate the response from the server. The filter can use a defined expression or the evaluator can use a defined string to help determine the server's health. If the evaluation passes, the server is healthy; otherwise, the health state of the server is down. The response must be valid XML. The response is analyzed with the XSL health check filter against the defined XPath expression.

The *timeout* argument specifies how much time can expire before an attempt to complete a health check fails. However, if the request hangs because the server does not respond, the health check timeout compares only the actual time that the request took. In this case, the health check timeout is not used to interrupt the transaction. The *enforce-timeout* argument overrides this behavior and forces the timeout to interrupt the transaction.

The *frequency* argument specifies the number of seconds between health checks. This frequency value is used sequentially such that the health check consecutively queries each load balancer member only after the prior health check completes. If a server hangs for a long time, all other health checks for the other members are delayed by that amount of time. The *independent-checks* argument overrides this behavior and makes each health check independent.

Examples

- Specify a periodic health check for members.

```
# health-check on cgi-bin/x.cgi 80 Standard
on store:///identity.xsl 4 60 / store:///healthcheck.xsl sslProxy1
#
```

- Specify a periodic health check for members of the `test1` load balancer group. The submode commands specify that the target URI is `cgi-bin/x.cgi` and the target port is `80`. The health check timeout value is used, and health checks run independently. All other properties use the default values. The health check is saved on exit.

```
# loadbalancer-group test1
# health-check
# admin-state on
# target-uri cgi-bin/x.cgi
# target-port 80
# enforce-timeout on
# independent-checks on
# exit
```

Related reference

[“server” on page 854](#)

This command adds a member to the load balancer group.

masquerade

This command specifies the host name to provide to the remote server.

Syntax

```
masquerade { on | off }
```

Parameters

on

Passes the name of the load balanced group to the remote server.

off

Passes the name of the member server to the remote server. This setting is the default value.

Example

Pass the name of the load balancer group as the host name to the remote server.

```
# masquerade on
#
```

server

This command adds a member to the load balancer group.

Syntax

```
server address [weight] [mapped-port] ["" ] [health-port]
```

Parameters

address

Specifies the name or IP address of the server.

weight

For weighted algorithms: Specifies the relative weight (preference). Enter a value in the range 1 - 65000. The default value is 1.

mapped-port

Specifies the port on the real server. If nonzero, the associated real server is contacted on this port. Normally the real server is contacted on the same port number as the one for the virtual server. In this case, retain the default value of 0. If services run on different ports for different members of the group, define this value.

This port is used for IMSConnect health check.

health-port

Specifies the port to test. Retain the default value of 0 to use the port that is defined for this load balancer group.

Guidelines

The **server** command adds a server to the load balancer group. When you define a member server, optionally specify the ports to use for sending client requests and for the health check.

- For the first-alive algorithm, the sequence is significant. The first server is the primary server, while subsequent entries serve as backup servers. For all other algorithms, the sequence is not significant.
- For weighted algorithms, use the weight parameter to specify the relative preference of a server. The greater the value, the more likely this server is to receive a connection request. Assume that the load balancer group has the following members.
 - Member A with a weight of 100.
 - Member B with a weight of 60.
 - Member C with a weight of 40.

Because of the weights, member A receives 50% of requests, member B receives 30% of requests, and member C receives 20% of requests.

If you use a weighted algorithm, there is a limit of 64 servers per load balancer group. For all other algorithms, the limit is 512 servers per load balancer group.

Example

Add `ragnarok.appliance.com` with a weight of 5 to the load balancer group.

```
# server ragnarok.appliance.com 5  
#
```

Related reference

[“algorithm” on page 849](#)

This command specifies the server selection algorithm.

[“health-check” on page 851](#)

This command defines the periodic health check procedure.

try-every-server

This command specifies the retry-behavior for a failed attempt.

Syntax

try-every-server { on | off }

Parameters

on

Sends the requests to each server until one responds or all fail. Each server that fails is set to the `softdown` state.

off

Does not attempt to contact other members. This setting is the default value.

Guidelines

The **try-every-server** sends the request to each server until one responds or all fail. This command applies only when none of the group members are in the up state. Each server that fails is set to the `softdown` state.

Example

Contact all member server before failure.

```
# try-every-server on
#
```

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.

Command	Purpose
“active-timeout” on page 857	This command sets the timer to close an established and active connection to the server.
“admin-state” on page 708	This command sets the administrative state for the configuration.
“ansi-color” on page 857	This command manages a multicolored console log.
“archive-mode” on page 858	This command specifies an archival behavior for file-based logs.
“backup” on page 859	This command specifies a backup for the current log.
“connect-timeout” on page 859	This command sets the time to wait to establish a connection.
“email-address” on page 860	This command specifies the email address of a remote recipient of SMTP log messages.

Table 57. Log Target commands (continued)

Command	Purpose
“event” on page 860	This command adds an event class and a priority to the log target.
“event-code” on page 861	This command specifies an event code to add to the log target.
“event-detection” on page 861	This command suppresses identical events in the log target.
“event-filter” on page 862	This command specifies an event code to exclude from the log target.
“facility” on page 862	This command specifies the syslog facility.
“feedback-detection” on page 863	This command suppresses events from the logging subsystem itself.
“format” on page 863	This command specifies the format in which events are added to the log.
“idle-timeout” on page 864	This command sets the idle timer for the logging target.
“ip-filter” on page 864	This command adds an IP address to include in the log target.
“local-address” on page 865	This command specifies the local address over which log events are transmitted to a remote recipient.
“local-file” on page 865	This command specifies a local file for log messages.
“local-ident” on page 866	This command sets the string that the remote recipient uses to identify this log target.
“object” on page 867	This command adds an object filter to the current log target.
“priority” on page 867	This command sets the service-level priority for the log target.
“rate-limit” on page 868	This command sets the maximum rate of events to log.
“remote-address” on page 868	This command specifies the destination address of log messages or the log itself.
“remote-directory” on page 870	This command specifies the remote directory where uploaded logs are stored.
“remote-login” on page 871	This command specifies the user name to use to upload a log file to a remote server.
“remote-port” on page 871	This command specifies the listening port on the remote server.
“rotate” on page 872	This command sets the maximum number of file rotations.
“sender-address” on page 873	This command specifies the email address of the sender.
“size” on page 873	This command sets the maximum size of a local log file.
“smtp-domain” on page 874	This command specifies the domain name of the SMTP client.
“soap-version” on page 874	This command specifies the version of SOAP to use.
“ssl” on page 875	This command associates an SSL proxy profile for SOAP-based log over HTTPS.
“suppression-period” on page 876	This command sets the interval to suppress identical events.
“timestamp” on page 877	This command specifies the time stamp format.
“trigger” on page 877	This command creates an event trigger.
“type” on page 878	This command identifies the logging model.

Table 57. Log Target commands (continued)

Command	Purpose
“upload-method” on page 880	This command sets the protocol to upload a file-based log to a remote site.
“url” on page 881	This command sets the destination for SOAP-based log entries.

Related reference

[“logging target” on page 805](#)

This command enters Logging mode.

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.

Related reference

[“connect-timeout” on page 859](#)

This command sets the time to wait to establish a connection.

[“idle-timeout” on page 864](#)

This command sets the idle timer for the logging target.

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.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

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
```

Related reference

[“remote-address” on page 868](#)

This command specifies the destination address of log messages or the log itself.

[“remote-directory” on page 870](#)

This command specifies the remote directory where uploaded logs are stored.

[“remote-login” on page 871](#)

This command specifies the user name to use to upload a log file to a remote server.

[“type” on page 878](#)

This command identifies the logging model.

[“upload-method” on page 880](#)

This command sets the protocol to upload a file-based log to a remote site.

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.

Related reference

[“active-timeout” on page 857](#)

This command sets the timer to close an established and active connection to the server.

[“idle-timeout” on page 864](#)

This command sets the idle timer for the logging target.

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.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

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
```

Related reference

[“type” on page 878](#)

This command identifies the logging model.

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
#
```

Related reference

[“event-filter” on page 862](#)

This command specifies an event code to exclude from the log target.

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.

Related reference

[“suppression-period” on page 876](#)

This command sets the interval to suppress identical events.

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
#
```

Related reference

[“event-code” on page 861](#)

This command specifies an event code to add to the log target.

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
#
```

Related reference

[“local-address” on page 865](#)

This command specifies the local address over which log events are transmitted to a remote recipient.

[“local-ident” on page 866](#)

This command sets the string that the remote recipient uses to identify this log target.

[“remote-address” on page 868](#)

This command specifies the destination address of log messages or the log itself.

[“type” on page 878](#)

This command identifies the logging model.

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.

Related reference

[“show logging” on page 896](#)

This command displays a specified appliance log.

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.

Related reference

[“active-timeout” on page 857](#)

This command sets the timer to close an established and active connection to the server.

[“connect-timeout” on page 859](#)

This command sets the time to wait to establish a connection.

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
```

Related reference

[“type” on page 878](#)

This command identifies the logging model.

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.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

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.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

nfs-file

This command specifies the path to the NFS mount file.

Syntax

nfs-file *file*

Parameters

file

Specifies the path to the log file relative to the NFS mount point.

Guidelines

The file name can use only the characters a-z, A-Z, 0-9, or an underscore. The path can have subdirectories that are delimited by a slash.

The file must have write permission.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

nfs-static-mount

This command assigns a static mount.

Syntax

nfs-static-mount *name*

Parameters

name

Specifies the name of an NFS static mount.

Guidelines

When the log type is `nfs`, specify the NFS static mount point to write the log over NFS.

To create an NFS static mount, use the Global **`nfs-static-mount`** command.

Related reference

[“nfs-file” on page 866](#)

This command specifies the path to the NFS mount file.

[“type” on page 878](#)

This command identifies the logging model.

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
#
```

Related reference

“type” on page 878

This command identifies the logging model.

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
#
```

Related reference

[“archive-mode” on page 858](#)

This command specifies an archival behavior for file-based logs.

[“remote-port” on page 871](#)

This command specifies the listening port on the remote server.

[“type” on page 878](#)

This command identifies the logging model.

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 *%2Ffile-path*.

The path in the URL resolves to *//file-path* for SCP or SFTP and */%2Ffile-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/
#
```

Related reference

[“archive-mode” on page 858](#)

This command specifies an archival behavior for file-based logs.

[“remote-address” on page 868](#)

This command specifies the destination address of log messages or the log itself.

[“type” on page 878](#)

This command identifies the logging model.

[“upload-method” on page 880](#)

This command sets the protocol to upload a file-based log to a remote site.

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/
#
```

Related reference

[“archive-mode” on page 858](#)

This command specifies an archival behavior for file-based logs.

[“remote-address” on page 868](#)

This command specifies the destination address of log messages or the log itself.

[“remote-directory” on page 870](#)

This command specifies the remote directory where uploaded logs are stored.

[“type” on page 878](#)

This command identifies the logging model.

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
#
```

Related reference

[“remote-address” on page 868](#)

This command specifies the destination address of log messages or the log itself.

[“type” on page 878](#)

This command identifies the logging model.

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`.

Related reference

[“archive-mode” on page 858](#)

This command specifies an archival behavior for file-based logs.

[“type” on page 878](#)

This command identifies the logging model.

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`.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

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`.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

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
#
```

Related reference

[“type” on page 878](#)

This command identifies the logging model.

[“upload-method” on page 880](#)

This command sets the protocol to upload a file-based log to a remote site.

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.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

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.

Related reference

[“type” on page 878](#)

This command identifies the logging model.

[“url” on page 881](#)

This command sets the destination for SOAP-based log entries.

ssl-client

This command associates an SSL client profile for SOAP-based or syslog log events.

Syntax

ssl-client *name*

no ssl-client

Parameters

name

Specifies the name of an SSL client profile.

Guidelines

The **ssl-client** command specifies the SSL client profile to secure connections between the appliance and its targets.

To create an SSL client profile, use the Crypto **ssl-client** command. To remove the profile, use the **no ssl-client** command.

This command is relevant only when the following conditions are met:

- The log type set by the **type** command is **soap**, **syslog-tcp**, or **syslog-ng**.
- The type set by the **ssl-client-type** command is **client**.

Related reference

[“ssl-client” on page 747](#)

This command enters SSL client profile mode.

[“type” on page 878](#)

This command identifies the logging model.

ssl-client-type

This command sets the type of the SSL profile for SOAP-based or syslog log events.

Syntax

Use an SSL client profile

ssl-client-type *client*

Use an SSL proxy profile (deprecated)

ssl-client-type *proxy*

Parameters

proxy (deprecated)

Uses the SSL proxy profile with the cryptographic profiles to secure connections. This setting is the default value for backward compatibility.

client

Uses the SSL client profile to secure connections.

Guidelines

The `ssl-client-type` command sets the SSL profile type to secure connections between the DataPower® Gateway and its targets. You can use an SSL proxy profile or an SSL client profile.

- The SSL proxy profile is deprecated. Check whether your configuration uses an SSL proxy profile. If yes, modify the configuration to use an SSL client profile to secure connections. To specify an SSL proxy profile, use the **ssl** command.
- To specify an SSL client profile, use the **ssl-client** command.

This command is relevant only when the log type set by the `type` command is `soap`, `syslog-tcp`, or `syslog-ng`.

Related reference

[“ssl-client” on page 747](#)

This command enters SSL client profile mode.

[“type” on page 878](#)

This command identifies the logging model.

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.

Related reference

[“event-detection” on page 861](#)

This command suppresses identical events in the log target.

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.

Related reference

[“email-address” on page 860](#)

This command specifies the email address of a remote recipient of SMTP log messages.

[“facility” on page 862](#)

This command specifies the syslog facility.

[“local-address” on page 865](#)

This command specifies the local address over which log events are transmitted to a remote recipient.

[“local-file” on page 865](#)

This command specifies a local file for log messages.

[“local-ident” on page 866](#)

This command sets the string that the remote recipient uses to identify this log target.

[“rate-limit” on page 868](#)

This command sets the maximum rate of events to log.

[“remote-address” on page 868](#)

This command specifies the destination address of log messages or the log itself.

[“rotate” on page 872](#)

This command sets the maximum number of file rotations.

[“size” on page 873](#)

This command sets the maximum size of a local log file.

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/
#
```

Related reference

[“archive-mode” on page 858](#)

This command specifies an archival behavior for file-based logs.

[“type” on page 878](#)

This command identifies the logging model.

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.

This command must be run from the IBM MQ administration mode. If the system is in the IBM MQ administration mode the prompt includes `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 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.

Remote appliance(s) unavailable

This status means that the replication connection to the other appliance has been lost.

Partitioned

Queue manager data on the appliances is out of step, and cannot be automatically resolved.

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.

Inconsistent

The status is displayed on a secondary appliance during the initial synchronization of a queue manager if connection has been lost and synchronization was interrupted. The secondary appliance cannot provide high availability functionality until the initial synchronization has completed.

Resource failed

The physical interface associated with the floating IP used by the queue manager has been reconfigured. The queue manager is stopped and cannot be restarted until the original configuration is restored. This can also apply if an aggregated interface associated with the floating IP is changed.

This situation can also arise if the queue manager has been configured for disaster recovery and the eth20 interface has been reconfigured.

To recover, you must restore the original interface configuration, or remove the disaster recovery configuration. You can also delete the queue manager.

V 9.1.0 Migration Pending

When you have upgraded the appliance from Version 8.0 or Version 9.0 to Version 9.1, high availability queue managers display this status until you run the `sethagrp -u` command to upgrade them, and the other appliance is upgraded to Version 9.1 too. .

- 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 last synchronized. This information is shown when the status is `Partitioned` or `Inconsistent`.
- The time and date that the data was last in sync. This is displayed whenever data is out of sync.

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 an in-progress synchronization has been interrupted. 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 **crtldrprimary** command has been run, to specify that a queue manager has the primary role, but no **crtldrsecondary** 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 last synchronized. This information is shown only when the status is *Partitioned*, *Inconsistent*, or *Remote appliance(s) unavailable*.
- The time and date that the data was last in sync. This is displayed whenever data is out of sync.
- 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 (mqc1i)#`. To enter the IBM MQ administration mode, enter `mqc1i` 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 for 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:

Argument	Description
aliases	Lists the currently defined command macros.
cert-monitor	Displays the current cert-monitor configuration.
certificate	Shows current certificate aliases.
crypto-mode	
cryptoprofile	
dns	Displays the current DNS configuration, see “DNS commands” on page 764.
dns-cache	Shows the DNS hosts held in the DNS cache.
ethernet	Displays the current Ethernet configuration, see “Ethernet commands” on page 769.
failure-notification	
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 800.
link-aggregation	Displays the current link-aggregation configuration, see “Link aggregation commands” on page 837.
name-servers	Shows currently defined name servers.
netarp	
network	Displays the current network configuration, see “Network commands” on page 911.

Table 58. 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 750.
op-state	Displays information about all the configuration objects.
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 997.

show audit-log

This command displays the contents of the audit log.

Syntax

```
show audit-log [-np] [rotation { 1 | n }]
show audit-log [-np] user [rotation { 1 | n }]
show audit-log [-np] date [rotation { 1 | n }]
show audit-log [-np] time [rotation { 1 | n }]
show audit-log [-np] address [rotation { 1 | n }]
```

Parameters

-np

Indicates no pagination.

user

Sorts the events in the audit log alphabetically by user name.

address

Sorts the events in the audit log numerically by IP address.

date

Sorts the events in the audit log numerically by date.

time

Sorts the events in the audit log numerically by time.

rotation {1 | n}

Specify the audit log rotation to show. Without this parameter, shows the audit-log.1, audit-log.2, and so on).

Guidelines

The **show audit-log** command displays the audit log with or without pagination. Use the **-np** keyword to display the audit log without pagination. Use the **user**, **date**, **time**, or **address** keyword to indicate the sorting sequence. Use the **rotation** keyword to indicate the audit log to show (audit-log.1, audit-log.2, and so on).

The **date** and **time** keywords are equivalent.

Examples

- Display the events in the current audit log file (`audit:///audit-log`) in date order.

```
# show audit-log date
#
```

- Display the events in the second rotation of the audit log file (`audit:///audit-log.2`) by date without pagination.

```
show audit-log -np date rotation 2
```

show audit-search

This command searches the audit log and displays matching events.

Syntax

```
show audit-search [-np] user name
```

```
show audit-search [-np] date start [end]
```

```
show audit-search [-np] time start [end]
```

```
show audit-search [-np] address address[/netmask]
```

Parameters

-np

Indicates no pagination.

user *name*

Displays events in the audit log for the specified user.

date *start* [*end*]

Displays events in the audit log from the specified start date to optional end date. Without an end date, displays events to the most recent date.

time *start* [*end*]

Displays events in the audit log from the specified start time to the optional end time. Without an end time, displays events until 23:59:59.

address *address*[/*netmask*]

Displays events in the audit log for the specified IP address or, if you use a netmask, the IP address range.

Guidelines

The **show audit-search** command searches the audit log and displays events that match the specified criteria. Use the **-np** parameter to indicate no pagination for the output.

Examples

- Display events in the audit log for the `joesmith` account one screen at a time.

```
# show audit-search user joesmith
#
```

- Display events in the audit log from February 10, 2008 onward one screen at a time.

```
# show audit-search date 20080210
#
```

- Display events in the audit log from IP address `10.10.10.15` upward as one continuous list.

```
# show audit-search -np address 10.10.10.15
#
```

- Display events in the audit log from the IP address in the range of 10.10.10.0 through 10.10.10.255 one screen at a time.

```
# show audit-search address 10.10.10.0/24
#
```

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
```

Related reference

[“show time” on page 910](#)

This command displays the current time and appliance uptime.

M 2000 M 2001 **show fibre-channel-hba**

This command displays the state of the fibre channel adapters.

Syntax

show fibre-channel-hba

Guidelines

The **show fibre-channel-hba** command shows the current state of the fibre channel adapters on the appliance.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

Example

```
# show fibre-channel-hba
fibre-channel-hba: fch1 [up]
-----
admin-state enabled

fibre-channel-hba: fch2 [up]
-----
admin-state enabled
```

M 2000 M 2001 *show fibre-channel-hba-status*

This command displays the current status of the fibre channel adapters.

Syntax

show fibre-channel-hba-status

Guidelines

The **show fibre-channel-hba-status** command shows the current status of the fibre channel adapters on the appliance.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

Example

```
mqa# show fibre-channel-hba-status
```

HBA	Op-State	WWPN	Port	state	Port speed	Port type	Supported speeds
fch1	up	10:00:00:90:fa:8e:0a:c2	online	8 Gbit	nport	4 Gbit, 8 Gbit, 16 Gbit	
fch2	up	10:00:00:90:fa:8e:0a:c3	online	8 Gbit	nport	4 Gbit, 8 Gbit, 16 Gbit	

M 2000 M 2001 *show fibre-channel-luns*

This command displays the available LUNs that the appliance knows about.

Syntax

show fibre-channel-luns

Guidelines

The **show fibre-channel-luns** command shows the LUNs (identified by LUIDs) that are available to be used by volume definitions. Each volume must use a different LUN, so that each queue manager has a dedicated LUN.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

Example

```
# show fibre-channel-luns
```

LUID	HBA	SCSI LUN	WWPN
600507680181804D98000000000001B47	fch2	0	50:05:07:68:01:10:26:65
600507680181804D98000000000001B47	fch2	0	50:05:07:68:01:10:27:05
600507680181804D98000000000001B47	fch1	0	50:05:07:68:01:30:26:65
600507680181804D98000000000001B47	fch1	0	50:05:07:68:01:30:27:05
600507680181804D980000000000235C	fch2	1	50:05:07:68:01:10:26:65
600507680181804D980000000000235C	fch2	1	50:05:07:68:01:10:27:05
600507680181804D980000000000235C	fch1	1	50:05:07:68:01:30:26:65
600507680181804D980000000000235C	fch1	1	50:05:07:68:01:30:27:05

M2000 M2001 *show fibre-channel-volume*

This command displays the state of the fibre channel volumes.

Syntax

show fibre-channel-volume

Guidelines

The **show fibre-channel-volume** command shows the current state of the volumes defined on the appliance.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

Example

```
# show fibre-channel-volumefibre-channel-volume: jps2 [up] (new)
-----
admin-state enabled
lun-uid 600507680181804D980000000000235C
use-multipath on

fibre-channel-volume: mandyvol [up]
-----
admin-state enabled
lun-uid 600507680181804D9800000000001B47
use-multipath on
```

show fibre-channel-volume-status

This command displays the current status of the fibre channel volumes.

Syntax

show fibre-channel-volume-status

Guidelines

The **show fibre-channel-volume-status** command shows the current status of the volumes defined on the appliance.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

Example

```
mqa# show fibre-channel-volume-status

Volume  LUID                               OpStatus Link Status
-----  -
jps2    600507680181804D980000000000235C  up       up
mandyvol 600507680181804D9800000000001B47  up       up
```

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 information about the current firmware.

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.

Related reference

[“show firmware-version” on page 892](#)

This command displays the current firmware version, without image type and installation date.

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 level of the firmware.
- The build number of the firmware.
- The date on which the firmware was built.
- The delivery type of the firmware.
- 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 (DPOS).
- The type of XML accelerator.
- The IBM Machine Type of the appliance.
- The IBM Model Type of the appliance.

Related reference

[“show firmware” on page 891](#)

This command displays information about the current firmware.

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

Related concepts

[“IPMI user commands” on page 829](#)

You can use the IPMI user commands to create or modify an IPMI users on the IBM MQ Appliance.

[“IPMI LAN channel commands” on page 831](#)

You can use the IPMI LAN channel commands to create or modify an IPMI LAN channel on the IBM MQ Appliance.

show ipmi-user

This command shows the current IPMI user configuration.

Syntax

show ipmi-user

Related concepts

[“IPMI user commands” on page 829](#)

You can use the IPMI user commands to create or modify an IPMI users on the IBM MQ Appliance.

[“IPMI LAN channel commands” on page 831](#)

You can use the IPMI LAN channel commands to create or modify an IPMI LAN channel on the IBM MQ Appliance.

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.

Related reference

[“show link-aggregation-status” on page 894](#)

This command provides statistics for aggregate interfaces.

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.

Related reference

[“show link-aggregation-member-status” on page 893](#)

This command lists members in link aggregation interfaces.

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`, `ip6tn10`, `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.

Related reference

[“show logging” on page 896](#)

This command displays a specified appliance log.

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.

Related reference

[“show log” on page 896](#)

This command displays the appliance 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
```

M 2000 M 2001 **show multipath**

This command shows the routes available for volumes that are using multipath.

Syntax

show multipath

Guidelines

The **show multipath** command shows the routes available for each volume.

The following example shows the information displayed for the volume named 'volume1':

Volume name	Group number	Group priority	HBA	Bus	LUN	Group state	Path state	Device state
volume1	1	50	fch1	0	0	active	active	ready
volume1	1	50	fch2	0	0	active	active	ready
volume1	2	10	fch1	0	0	enabled	active	ready
volume1	2	10	fch2	0	0	enabled	active	ready

The **show multipath** command shows route information only for multipath volumes that are in OP-State UP. Additionally, if a route is lost to the SAN Storage, that route is no longer be listed in the table.

Note: This command is available only on M2000 and M2001 appliances, it is not available on M2002 appliances.

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, ip6tn10, 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, ip6tn10, 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 op-state

This command displays information about configuration options.

Syntax

show op-state

Guidelines

Displays information about all the configuration objects on the appliance.

Context

Available in Global Configuration mode.

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
  iyvush
```

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

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.

Example

The IBM MQ Appliance has four SSDs, configured as a RAID 10 array. The appliance configures two RAID 1 arrays as mirrors as part of the RAID 10 configuration, and the **show raid-array** shows the status of the two RAID 1 arrays. Entering the **show raid-array** command displays an output similar to the following example:

```
mqa# show raid-array
```

Controller ID	Array ID	Logical drive ID	RAID level	Number of physical drives	Total size
1	1	1	raidLevel1	2	3051757
1	2	1	raidLevel1	2	3051757

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 milliamperes-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

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.

Example

The IBM MQ Appliance M2002 has four SSDs, configured as a RAID 10 array. Entering the **show raid-logical-drive** command displays an output similar to the following example:

```
mqa# show raid-logical-drive

Controller ID Logical drive ID Logical drive name RAID level  Number of physical drives State
Initialization progress Read policy      Write policy Cache policy Access policy Bad block
exists Total size
-----
1          1          raid0          raidLevel10 4          readWrite          optimal
noInitInProgress          adaptiveReadAhead writeBack      directIO
noBadBlock          6103514
```

show raid-physical-drive

This command displays the status of the RAID physical drive.

Syntax

show raid-physical-drive

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.

Example

The IBM MQ Appliance M2002 has four SSDs, configured as a RAID 10 array. Entering the **show raid-physical-drive** command displays an output similar to the following example:

```
mqa# show raid-physical-drive

Controller ID Device ID Array ID Logical drive ID Logical drive name Position State Progress
Raw size Coerced size Drive type Interface speed SAS address Vendor ID
Product ID Revision Specific information Failure Temperature
-----
1 8 1 1 raid0 B7S2 online 0
3052360 3051757 sas notApplicable 5000c5003028294900000000000000 IBM-ESXS
ST3200FM0023 4C19 ZAM159MR4C194C194C19 No 35 / 95
1 9 1 1 raid0 B5S1 online 0
3052360 3051757 sas notApplicable 5000c5003028296100000000000000 IBM-ESXS
ST3200FM0023 4C19 ZAM159GL4C194C194C19 No 33 / 91
1 10 2 1 raid0 B5S2 online 0
3052360 3051757 sas notApplicable 5000c5003028289500000000000000 IBM-ESXS
ST3200FM0023 4C19 ZAM159AJ4C194C194C19 No 33 / 91
1 11 2 1 raid0 B7S1 online 0
3052360 3051757 sas notApplicable 5000c5003028284100000000000000 IBM-ESXS
ST3200FM0023 4C19 ZAM1599W4C194C194C19 No 35 / 95
```

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 displayed:

- The disk number. The value is 1 or 2 (for M2001 appliances), or 1, 2, 3, or 4 (for M2002 appliances).
- 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 ZAM11500 86 100
```

Related reference

[“show raid-physical-drive” on page 903](#)

This command displays the status of the RAID physical drive.

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.

Related reference

“show tcp-table” on page 909

This command lists the current TCP connections.

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.

Related reference

“show tcp-connections” on page 909

This command lists the number of TCP connections in specific states.

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.

Related reference

[“show tcp-connections” on page 909](#)

This command lists the number of TCP connections in specific states.

[“show tcp-table” on page 909](#)

This command lists the current TCP connections.

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
```

Related reference

[“show clock” on page 888](#)

This command displays the current time and appliance uptime.

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.

Related reference

[“show firmware-version” on page 892](#)

This command displays the current firmware version, without image type and installation date.

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
```

Related reference

[“arp-retries” on page 912](#)

This command sets the number of times the appliance attempts a failed ARP request.

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
```

Related reference

[“arp-interval” on page 911](#)

This command sets the time interval between ARP attempts.

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
```

Related reference

[“relax-interface-isolation” on page 915](#)

This command controls whether to allow packets from a wrong interface when both interfaces are on the same subnet.

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
```

Related reference

[“relax-interface-isolation” on page 915](#)

This command controls whether to allow packets from a wrong interface when both interfaces are on the same subnet.

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
```

Related reference

[“destination-routing” on page 912](#)

This command controls how the appliance determines the route to return the response.

[“disable-interface-isolation” on page 913](#)

This command controls interface isolation.

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.

Related reference

[“arp-interval” on page 911](#)

This command sets the time interval between ARP attempts.

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.

RBM setting commands

Use the RBM commands to configure role based management on the IBM MQ Appliance.

You enter the commands in RBM configuration mode:

1. From the appliance command line, enter global configuration mode:

```
config
```

2. From global configuration mode, type `rbm` to enter RBM configuration mode.
3. Type `exit` to leave the configuration mode and save your changes, then type `exit` again to leave global configuration mode.

au-cache-mode

This command sets the caching mode for authentication results.

Syntax

au-cache-mode { absolute | disabled | maximum | minimum }

Parameters

absolute

Caches the results of user authentications for a period of time specified by the **au-cache-ttl** command (the explicit time-to-live).

disabled

Disables caching. The system will not cache any results and instead always authenticate every time a user requests access.

maximum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the maximum of the two values. This setting is the default value.

minimum

Compares the explicit TTL to the TTL contained in the response (if any) and cache authentication results for the minimum of the two values.

Guidelines

The **au-cache-mode** command establishes the required caching mode for authentication results. Use the **au-cache-ttl** command to establish the explicit TTL.

Example

Cache authentication results for the maximum amount of time.

```
# au-cache-mode maximum
```

Related reference

[“au-cache-ttl” on page 918](#)

This command specifies the time-to-live for cached authentication results.

au-cache-ttl

This command specifies the time-to-live for cached authentication results.

Syntax

au-cache-ttl *seconds*

Parameters***seconds***

Specifies the time-to-live (TTL) in seconds. Enter a value in the range 1 - 86400. The default value is 600.

Guidelines

The **au-cache-ttl** command defines the explicit TTL in seconds for cached authentication results. This value is compared against the TTL in the authentication response in accordance with the cache mode, as defined with the **au-cache-mode** command.

Example

Set the TTL to 5 minutes.

```
# au-cache-ttl 300
```

Related reference

[“au-cache-mode” on page 917](#)

This command sets the caching mode for authentication results.

au-info-url

This command specifies the URL of the authentication XML file.

Syntax

au-info-url *URL*

Parameters

URL

Specifies the location of the XML file.

Guidelines

The **au-info-url** command defines the fully qualified file name (URL) of the XML file for authentication. This command is relevant when the authentication method, as defined with the **au-method** command, is `xmlfile`.

Example

Identify the `RBM-AU.xml` file in the `local:` directory as the authentication XML file.

```
# au-method xmlfile
# au-info-url local:///RBM-AU.xml
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[au-ldap-bind-dn](#)

This command specifies the login DN to access an LDAP server.

Syntax

au-ldap-bind-dn *DN*

Parameters

DN

Specifies the login DN (distinguished name).

Guidelines

The **au-ldap-bind-dn** command specifies the login DN to access the target LDAP server. This command is relevant when the authentication method, as defined with the **au-method** command, is `ldap` and when the LDAP search for group name property, as defined with the **au-ldap-search** command, is enabled.

Beyond specifying the login DN to search the LDAP for the group name, you must use the **au-ldap-bind-password** command and optionally use the **au-ldap-parameters** command.

- The **au-ldap-bind-password** command specifies the user's password.
- The **au-ldap-parameters** command associates LDAP search parameters.

Example

Identify LDAP authentication with optional retrieval of the group DN.

```
# au-method ldap
# au-server-host ldap-1
# au-server-port 389
# au-ldap-search on
# au-ldap-bind-dn proxyuser
# au-ldap-bind-password p@Ssw0rd
```

Related reference

[“au-ldap-bind-password” on page 920](#)

This command specifies the password to access an LDAP server.

[“au-ldap-parameters” on page 921](#)

This command assigns LDAP search parameters for an LDAP search.

[“au-ldap-search” on page 922](#)

This command indicates whether to retrieve the DN with an LDAP search.

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

au-ldap-bind-password

This command specifies the password to access an LDAP server.

Syntax

au-ldap-bind-password *password*

Parameters

password

Specifies the password for the login DN.

Guidelines

The **au-ldap-bind-password** command specifies the password for the login DN to access the target LDAP server. This command is relevant when the authentication method, as defined with the **au-method** command, is ldap and when the LDAP search for group name property, as defined with the **au-ldap-search** command, is enabled.

Beyond specifying the login password to search the LDAP for the group name, you must use the **au-ldap-bind-dn** command and optionally use the **au-ldap-parameters** command.

- The **au-ldap-bind-dn** command specifies the user.
- The **au-ldap-parameters** command associates LDAP search parameters.

Example

Identify LDAP authentication with optional retrieval of the group DN.

```
# au-method ldap
# au-server-host ldap-1
# au-server-port 389
# au-ldap-search on
# au-ldap-bind-dn proxyuser
# au-ldap-bind-password p@Ssw0rd
```

Related reference

[“au-ldap-bind-dn” on page 919](#)

This command specifies the login DN to access an LDAP server.

[“au-ldap-parameters” on page 921](#)

This command assigns LDAP search parameters for an LDAP search.

[“au-ldap-search” on page 922](#)

This command indicates whether to retrieve the DN with an LDAP search.

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

au-ldap-parameters

This command assigns LDAP search parameters for an LDAP search.

Syntax

au-ldap-parameters *name*

Parameters

name

Specifies the name of the LDAP search parameters.

Guidelines

The **au-ldap-parameters** command assigns LDAP search parameters for an LDAP search. The search retrieves the user's distinguished name (DN).

This command is relevant only when LDAP search is enabled with the **au-ldap-search** command and when the authentication method is `ldap`, as defined with the **au-method** command.

Example

Identify LDAP authentication with optional retrieval of the group DN.

```
# au-method ldap
# au-server-host ldap-1
# au-server-port 389
# au-ldap-search on
# au-ldap-bind-dn proxyuser
# au-ldap-bind-password p@Ssw0rd
```

Related reference

[“au-ldap-search” on page 922](#)

This command indicates whether to retrieve the DN with an LDAP search.

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

au-ldap-readtimeout

This command sets the time that RBM authentication waits for a response from the LDAP server.

Syntax

au-ldap-readtimeout *seconds*

Parameters

seconds

Indicates the number of seconds to wait for a response from the LDAP server before the appliance closes the connection. Enter a value in the range 0 - 86400. The default value is 60. A value of 0 indicates that the connection never times out.

Guidelines

The **ldap-readtimeout** command specifies the number of seconds that the appliance waits for a response from the LDAP server before the appliance closes the LDAP connection. This command is relevant only when the authentication method, as defined with the **au-method** command, is `ldap`.

Specify a maximum time length of 60 seconds that the appliance waits for the response from the LDAP server.

```
# au-ldap-readtimeout 60
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

au-ldap-search

This command indicates whether to retrieve the DN with an LDAP search.

Syntax

au-ldap-search { on | off }

Parameters

on

Enables an LDAP search for the user's distinguished name (DN). The login name and LDAP search parameters are used as part of an LDAP search to retrieve the user's DN.

off

Disables an LDAP search for the user's DN. The login name with the LDAP prefix and suffix are used to construct the user's DN. This setting is the default value.

Guidelines

The **au-ldap-search** command indicates whether to retrieve the distinguished name with an LDAP search.

- When enabled, use the following command to complete the configuration.
 - The **au-ldap-bind-dn** command to specify the user's DN.
 - The **au-ldap-bind-password** command to specify the user's password.
 - The **au-ldap-parameters** command to associate an LDAP search parameters configuration.
- When disabled, use the following command to complete the configuration. The provided prefix and suffix form the DN to submit to the LDAP server.
 - The **ldap-prefix** command to specify the LDAP prefix to add to the user name.
 - The **ldap-suffix** command to specify the LDAP suffix to append to the user name.

This command is relevant when the authentication method, as defined with the **au-method** command, is ldap.

Example

Identify LDAP authentication with optional retrieval of the group DN.

```
# au-method ldap
# au-server-host ldap-1
# au-server-port 389
# au-ldap-search on
# au-ldap-bind-dn proxyuser
# au-ldap-bind-password p@Ssw0rd
```

Related reference

[“au-ldap-bind-dn” on page 919](#)

This command specifies the login DN to access an LDAP server.

[“au-ldap-bind-password” on page 920](#)

This command specifies the password to access an LDAP server.

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[“au-ldap-parameters” on page 921](#)

This command assigns LDAP search parameters for an LDAP search.

[“ldap-prefix” on page 927](#)

This command specifies the LDAP prefix to add to the user name to form the DN.

[“ldap-suffix” on page 929](#)

This command specifies the LDAP suffix to add to the user name to form the DN for RBM authentication.

au-method

This command specifies the RBM authentication method.

Syntax

au-method *method*

Parameters

method

Sets the RBM authentication method. The default value is `local`.

client-ssl

Uses a SSL certificate from a connection peer. Requires validation credentials associated with the **au-valcred** command.

ldap

Uses an LDAP server. Requires information about the LDAP server with the **au-server-host** and **au-server-port** commands.

local

Uses the user configuration that is maintained on the local system. Does not access external resources.

XML file

Uses a locally stored RBM Info file. Requires the location of the file with the **au-info-url** command.

Guidelines

The **au-method** command sets the authentication method for RBM. The selected method must be fully configured before invoking this command.

If the `admin` account is not configured with all permissions, the `admin` account is locked out of the GUI. Use the CLI to change this circumstance.

Related reference

[“au-info-url” on page 918](#)

This command specifies the URL of the authentication XML file.

[“au-server-host” on page 924](#)

This command specifies the IP address or domain name of a remote authentication server.

[“au-server-port” on page 924](#)

This command specifies the port on the remote authentication server.

[“au-valcred” on page 925](#)

This command assigns validation credentials for SSL client certificate authentication.

au-server-host

This command specifies the IP address or domain name of a remote authentication server.

Syntax

au-server-host *host*

Parameters

host

Specifies the IP address or domain name of the server.

Guidelines

The **au-server-host** command specifies the IP address or domain name of the authentication server.

When the authentication method is ldap, as defined with the **au-method** command, define the LDAP server in one of the following ways:

- The **au-server-host** and **au-server-port** commands
- The **loadbalancer-group** command

Example

Identify LDAP authentication with optional retrieval of the group DN.

```
# au-method ldap
# au-server-host ldap-1
# au-server-port 389
# au-ldap-search on
# au-ldap-bind-dn proxyuser
# au-ldap-bind-password p@Ssw0rd
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[“au-server-port” on page 924](#)

This command specifies the port on the remote authentication server.

[“loadbalancer-group” on page 930](#)

This command associates the load balancer group for RBM LDAP authentication.

au-server-port

This command specifies the port on the remote authentication server.

Syntax

au-server-port *port*

Parameters

port

Specifies the port number of the authentication server.

Guidelines

The **au-server-port** command specifies the listening port of the authentication server that is defined with the **au-server-host** command.

When the authentication method is ldap, as defined with the **au-method** command, define the LDAP server in one of the following ways:

- The **au-server-host** and **au-server-port** commands
- The **loadbalancer-group** command

Example

Identify LDAP authentication with optional retrieval of the group DN.

```
# au-method ldap
# au-server-host ldap-1
# au-server-port 389
# au-ldap-search on
# au-ldap-bind-dn proxyuser
# au-ldap-bind-password p@Ssw0rd
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[“au-server-host” on page 924](#)

This command specifies the IP address or domain name of a remote authentication server.

[“loadbalancer-group” on page 930](#)

This command associates the load balancer group for RBM LDAP authentication.

au-valcred

This command assigns validation credentials for SSL client certificate authentication.

Syntax

au-valcred *name*

Parameters

name

Specifies the name of the validations credentials.

Guidelines

The **au-valcred** command associates validations credentials to validate the identity in a client certificate from an SSL peer. This command is relevant when the authentication method, as defined with the **au-method** command, is `client-ssl`.

Use the Crypto **valcred** command to create validations credentials.

Example

Assign the `valCred-1` validations credentials for SSL client certificate authentication.

```
# au-method client-ssl
# au-valcred valCred-1
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

cli-timeout

This command specifies the time before the CLI session is closed because of inactivity.

Syntax

cli-timeout *seconds*

Parameters

seconds

Specifies the timeout value of the idle session in seconds. Enter a value in the range 0 - 65535. The default value is 0, which disables the timer.

Guidelines

The **cli-timeout** command specifies the amount of idle time in seconds before the CLI session is closed because of inactivity. When the session times out, you must reestablish a session and reauthenticate.

This command manages the session timeout for the CLI. The **idle-timeout** command in Web Management Service mode controls the session timeout for the GUI.

fallback-login

This command specifies whether to use local users if the primary authentication method fails.

Syntax

fallback-login { disabled | local | restricted }

Parameters

disabled

Indicates that no locally defined user can log on. This setting is the default value.

local

Indicates that all locally defined users can log on.

restricted

Indicates that only specific locally defined users can log on.

Guidelines

The **fallback-login** command indicates whether to use local user accounts as fallback users when the primary authentication method fails. With fallback users, locally defined users can log on to the appliance if the authentication method fails or in the event of a network outage that affects the primary authentication.

To limit fallback users to a specific set, use the `restricted` keyword. In this case, use the **fallback-user** command to define the specific, locally defined users to allow as fallback users.

The **fallback-login** command is relevant only when remote authentication. In other words, this command is relevant when the setting for the **au-method** is any value except `local`.

Note: On XI50z, this option is `local`. No administrator can modify this setting.

Examples

- Allow all locally defined users to log on.

```
# fallback-login local
```

- Designate bobsmith and joselopez as fallback users.

```
# fallback-login restricted
# fallback-user bobsmith
# fallback-user joselopez
```

- Disallow all locally defined users from logging on.

```
# fallback-login disabled
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[“fallback-user” on page 927](#)

This command adds a locally defined user as a fallback user.

fallback-user

This command adds a locally defined user as a fallback user.

Syntax

fallback-user *user*

no fallback-user *user*

Parameters

user

Specifies the name of a locally defined user.

Guidelines

The **fallback-user** command allows a locally defined user to be a fallback user. Run the **fallback-user** command for each fallback user.

This command is relevant when the **fallback-login** command is set to `restricted`.

Use the **no fallback-user** command to remove a user from the list of fallback users.

Example

Designate bobsmith and joselopez as fallback users.

```
# fallback-login restricted
# fallback-user bobsmith
# fallback-user joselopez
```

Related reference

[“fallback-login” on page 926](#)

This command specifies whether to use local users if the primary authentication method fails.

ldap-prefix

This command specifies the LDAP prefix to add to the user name to form the DN.

Syntax

ldap-prefix *prefix*

Parameters

prefix

Specifies an LDAP prefix.

Guidelines

The **ldap-prefix** command specifies the string to add as a prefix to the user name to form the distinguished name (DN) for LDAP authentication. The LDAP prefix and the user name are separated with a comma, and both are included within quotation marks.

If the LDAP prefix is `cn=` and the user name is `Bob Smith`, the beginning portion of the DN is `cn=Bob Smith`.

This command is relevant only when the **au-ldap-search** command is `off`.

Example

Set the LDAP prefix to `cn=`.

```
# ldap-prefix "cn="
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[“au-ldap-search” on page 922](#)

This command indicates whether to retrieve the DN with an LDAP search.

[“ldap-suffix” on page 929](#)

This command specifies the LDAP suffix to add to the user name to form the DN for RBM authentication.

ldap-sslproxy

This command associates the SSL proxy profile for LDAP authentication.

Syntax

ldap-sslproxy *name*

Parameters

name

Specifies the name of an SSL proxy profile.

Guidelines

The **ldap-sslproxy** command associates the SSL proxy profile to secure communication with the LDAP server during LDAP authentication. When specified, LDAP communication uses the configuration in this SSL proxy profile. If not specified, the communication is nonsecure.

To create an SSL proxy profile, use the Global **sslproxy** command.

This command is relevant only when the authentication method, as specified with the **au-method** command, is `ldap`.

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

ldap-suffix

This command specifies the LDAP suffix to add to the user name to form the DN for RBM authentication.

Syntax

ldap-suffix *suffix*

Parameters

suffix

Specifies an LDAP suffix.

Guidelines

The **ldap-suffix** command specifies the string to add after the user name to form the base distinguished name (DN) for LDAP authentication. The LDAP suffix and the user name are separated with a comma, and both are included within quotation marks.

For example, if LDAP suffix is `O=example.com` and the user name is Bob, the DN is `CN=Bob,O=example.com`.

This command is relevant only in the following situation:

- The **au-method** command is set to `ldap`.
- The **au-ldap-search** command is set to `off`.

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[“au-ldap-search” on page 922](#)

This command indicates whether to retrieve the DN with an LDAP search.

[“ldap-prefix” on page 927](#)

This command specifies the LDAP prefix to add to the user name to form the DN.

ldap-version

This command specifies the LDAP version to access the LDAP server for RBM authentication.

Syntax

ldap-version { v2 | v3 }

Parameters

v2

Uses LDAP version 2 as the protocol. This setting is the default value.

v3

Uses LDAP version 3 as the protocol.

Guidelines

The **ldap-version** command specifies the LDAP version for RBM authentication. This command is relevant only when the **au-method** command is set to `ldap`.

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

loadbalancer-group

This command associates the load balancer group for RBM LDAP authentication.

Syntax

loadbalancer-group *name*

Parameters

name

Specifies the name of a load balancer group.

Guidelines

The **loadbalancer-group** command associates a load balancer group for LDAP authentication. When the authentication method is ldap, as defined with the **au-method** command, you must define the LDAP server in one of the following ways.

- The **au-server-host** and **au-server-port** commands
- The **loadbalancer-group** command

To create a load balancer group, use the Global **loadbalancer-group** command.

Example

Set the LDAP load balancer to LBGroup1.

```
# au-method ldap
# loadbalancer-group LBGroup1
# au-ldap-search on
# au-ldap-bind-dn proxyuser
# au-ldap-bind-password p@Ssw0rd
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[“au-server-host” on page 924](#)

This command specifies the IP address or domain name of a remote authentication server.

[“au-server-port” on page 924](#)

This command specifies the port on the remote authentication server.

lockout-duration

This command specifies the duration to lock out local accounts.

Syntax

lockout-duration *minutes*

Parameters

minutes

Specifies the number of minutes to lock out an account after the maximum number of failed login attempts is exceeded. A value of 0 indicates that accounts are locked out until reset by a privileged administrator. Enter a value in the range 0 - 1000. The default value is 1.

Guidelines

The **lockout-duration** command specifies the duration to lock out accounts after the maximum number of failed login attempts is exceeded. Define the maximum number of failed login attempts with the **max-login-failure** command. Instead of locking out an account for a specific duration, the account can be locked out until re-enabled by a privileged administrator. To lock out accounts until reset, set the duration to 0.

Note: The **lockout-duration** command applies to all local accounts, which include the `admin` account. When the duration is 0, the `admin` account is locked out for 120 minutes or until reenabled by another administrator.

Examples

Enable lockout behavior for accounts that on the fifth login failure, the account is locked out until reset by a privileged administrator:

```
# lockout-duration 0
# max-login-failure 4
```

Related reference

[“max-login-failure” on page 931](#)

This command specifies whether to lock out a local user account after a specific number of failed login attempts.

max-login-failure

This command specifies whether to lock out a local user account after a specific number of failed login attempts.

Syntax

max-login-failure *count*

Parameters

count

Specifies the maximum number of failed login attempts to allow before lockout. A value of 0 disables account lockout. Enter a value in the range 0 64. The default value is 3.

Guidelines

The **max-login-failure** command defines the number of failed login attempts to allow before a successful login. If the value is 3 and the user failed three consecutive login attempts, the behavior on the next login attempt for this user is as follows:

- If failure, the account is locked out. The duration of the lockout depends on the value that is defined by the **lockout-duration** command.
- If successful, the account is not locked out and the count is reset.

Note: The **max-login-failure** command applies to all local accounts, which include the `admin` account. When the duration is 0, the `admin` account is locked out for 120 minutes or until reenabled by another administrator.

Examples

- Enable lockout behavior for accounts that on the fifth login failure, the account is locked out until reset by a privileged administrator:

```
# lockout-duration 0
# max-login-failure 4
```

- Disable lockout.

```
# max-login failure 0
```

Related reference

[“lockout-duration” on page 930](#)

This command specifies the duration to lock out local accounts.

mc-custom-url

This command specifies the URL of the RBM credential-mapping custom style sheet.

Syntax

mc-custom-url *URL*

Parameters

URL

Specifies the location of the style sheet.

Guidelines

The **mc-custom-url** command defines the fully qualified file name (URL) of the custom style sheet to map credentials. This command is relevant when the mapping credentials method, as defined with the **mc-method** command, is custom.

Example

Identify the RBM-MC.xml style sheet in the mapCred directory of the myserver.example.com server as the style sheet to map credentials. This file is retrieved over HTTPS.

```
# mc-method custom
# mc-custom-url https://myserver.example.com/mapCred/RBM.xml
```

Related reference

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

mc-info-url

This command specifies the URL of the mapping credentials XML file.

Syntax

mc-info-url *URL*

Parameters

URL

Specifies the location of the XML file.

Guidelines

The **mc-info-url** command defines the fully-qualified file name (URL) of the XML file for credentials mapping. This command is relevant when the mapping credentials method, as defined with the **mc-method** command, is xmlfile.

Example

Identify the RBM-MC.xml file in the local: directory.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
```

Related reference

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

mc-ldap-bind-password

This command specifies the password for the login DN to access an LDAP server.

Syntax

mc-ldap-bind-password *password*

Parameters

password

Specifies the password for the login DN.

Guidelines

The **mc-ldap-bind-password** command specifies the password for the login DN to access the target LDAP server.

This command is relevant only in the following situation:

- LDAP search is enabled with the **mc-ldap-search** command
- When the credentials mapping method, as defined with the **mc-method** command, is local or xmlfile.

Beyond the password for the login DN to search the LDAP for the group name, use the following commands to complete the configuration.

- How to connect to the LDAP server. Use either of the following approaches.
 - The **mc-server-host** and **mc-server-port** commands
 - The **mc-loadbalancer-group** command
- Optionally associate an SSL proxy profile with the **mc-ldap-sslproxy** command to secure communication with the LDAP server
- Specify the login DN to access the LDAP server with the **mc-ldap-bind-dn** command
- Optionally associate LDAP search parameters with the **mc-ldap-parameters** command

Example

Use a local XML file to map credentials and search the LDAP retrieve the distinguished name.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
# mc-ldap-search on
# mc-server-host ldap.mydomain.com
# mc-server-port 389
# mc-ldap-bind-dn "cn=proxyuser"
# mc-ldap-bind-password p@Ssw0rd
# mc-ldap-parameters ldap1-MC
```

Related reference

[“mc-ldap-bind-dn” on page 939](#)

This command specifies the login DN to access an LDAP server.

[“mc-ldap-parameters” on page 934](#)

This command assigns the LDAP search parameters to search LDAP.

[“mc-ldap-search” on page 936](#)

This command indicates whether to retrieve the group names with an LDAP search.

[“mc-ldap-sslproxy” on page 937](#)

This command assigns an SSL proxy profile with the LDAP credentials server.

[“mc-loadbalancer-group” on page 938](#)

This command assigns a load balancer group as the target for an LDAP search.

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

[“mc-server-host” on page 941](#)

This command specifies the host of the LDAP server.

[“mc-server-port” on page 943](#)

This command specifies the listening port on the LDAP server.

mc-ldap-parameters

This command assigns the LDAP search parameters to search LDAP.

Syntax

mc-ldap-parameters *name*

Parameters

name

Specifies the name of the LDAP search parameters.

Guidelines

The **mc-ldap-parameters** command assigns the LDAP search parameters configuration to perform an LDAP search. The search retrieves the user's group.

This command is relevant only in the following situation:

- LDAP search is enabled with the **mc-ldap-search** command
- When the credentials mapping method, as defined with the **mc-method** command, is `local` or `xmlfile`.

Beyond LDAP search parameters to search LDAP for the group name, use the following commands to complete the configuration.

- How to connect to the LDAP server. Use either of the following approaches.
 - The **mc-server-host** and **mc-server-port** commands
 - The **mc-loadbalancer-group** command
- Optionally associate an SSL proxy profile with the **mc-ldap-sslproxy** command to secure communication with the LDAP server
- Specify the login DN to access the LDAP server with the **mc-ldap-bind-dn** command
- Specify the user's password with the **mc-ldap-bind-password** command

Example

Use a local XML file to map credentials and search the LDAP retrieve the distinguished name.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
```

```
# mc-ldap-search on
# mc-server-host ldap.mydomain.com
# mc-server-port 389
# mc-ldap-bind-dn "cn=proxyuser"
# mc-ldap-bind-password p@Ssw0rd
# mc-ldap-parameters ldap1-MC
```

Related reference

[“mc-ldap-bind-dn” on page 939](#)

This command specifies the login DN to access an LDAP server.

[“mc-ldap-bind-password” on page 933](#)

This command specifies the password for the login DN to access an LDAP server.

[“mc-ldap-search” on page 936](#)

This command indicates whether to retrieve the group names with an LDAP search.

[“mc-ldap-sslproxy” on page 937](#)

This command assigns an SSL proxy profile with the LDAP credentials server.

[“mc-loadbalancer-group” on page 938](#)

This command assigns a load balancer group as the target for an LDAP search.

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

[“mc-server-host” on page 941](#)

This command specifies the host of the LDAP server.

[“mc-server-port” on page 943](#)

This command specifies the listening port on the LDAP server.

mc-ldap-readtimeout

This command sets the time that RBM credential mapping waits for a response from the LDAP server.

Syntax

mc-ldap-readtimeout *seconds*

Parameters

seconds

Indicates the number of seconds to wait for a response from the LDAP server before the appliance closes the connection. Enter a value in the range 0 - 86400. The default value is 60. A value of 0 indicates that the connection never times out.

Guidelines

The **mc-ldap-readtimeout** command specifies the number of seconds that the appliance waits for a response from the LDAP server before the appliance closes the LDAP connection. This command is relevant only when the credentials mapping method, as defined with the **mc-method** command, is `local` or `xmlfile`.

Example

Specify a maximum time length of 60 seconds that the appliance waits for the response from the LDAP server.

```
# mc-ldap-readtimeout 60
```

Related reference

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

mc-ldap-search

This command indicates whether to retrieve the group names with an LDAP search.

Syntax

mc-ldap-search { on | off }

Parameters

on

Enables an LDAP search for the user's group. The authenticated DN of the user with the LDAP search parameters are used as part of an LDAP search to retrieve the user's group.

off

Disables an LDAP search for the user's group. The authenticated identity of the user (DN or user group of local user) is used directly as the input credential. This setting is the default value.

Guidelines

The **mc-ldap-search** command indicates whether to retrieve the distinguished name with an LDAP search.

This command is relevant when the credentials mapping method, as defined with the **mc-method** command, is `local` or `xmlfile`.

When enabled, use the following commands to complete the configuration.

- How to connect to the LDAP server. Use either of the following approaches.
 - The **mc-server-host** and **mc-server-port** commands
 - The **mc-loadbalancer-group** command
- Optionally associate an SSL proxy profile with the **mc-ldap-sslproxy** command to secure communication with the LDAP server
- Specify the login DN to access the LDAP server with the **mc-ldap-bind-dn** command
- Specify the user's password with the **mc-ldap-bind-password** command
- Optionally associate LDAP search parameters with the **mc-ldap-parameters** command

Example

Use a local XML file to map credentials and search the LDAP retrieve the distinguished name.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
# mc-ldap-search on
# mc-server-host ldap.mydomain.com
# mc-server-port 389
# mc-ldap-bind-dn "cn=proxyuser"
# mc-ldap-bind-password p@Ssw0rd
# mc-ldap-parameters ldap1-MC
```

Related reference

[“mc-ldap-bind-dn” on page 939](#)

This command specifies the login DN to access an LDAP server.

[“mc-ldap-bind-password” on page 933](#)

This command specifies the password for the login DN to access an LDAP server.

[“mc-ldap-parameters” on page 934](#)

This command assigns the LDAP search parameters to search LDAP.

[“mc-ldap-sslproxy” on page 937](#)

This command assigns an SSL proxy profile with the LDAP credentials server.

[“mc-loadbalancer-group” on page 938](#)

This command assigns a load balancer group as the target for an LDAP search.

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

[“mc-server-host” on page 941](#)

This command specifies the host of the LDAP server.

[“mc-server-port” on page 943](#)

This command specifies the listening port on the LDAP server.

mc-ldap-sslproxy

This command assigns an SSL proxy profile with the LDAP credentials server.

Syntax

mc-ldap-sslproxy *name*

Parameters

name

Specifies the name of an SSL proxy profile.

Guidelines

The **mc-ldap-sslproxy** command assigns an SSL proxy profile to secure communication with the LDAP credentials server. When specified, LDAP communication uses the configuration in the assigned SSL proxy profile. If not specified, communications is nonsecure.

This command is relevant only in the following situation:

- LDAP search is enabled with the **mc-ldap-search** command
- When the credentials mapping method, as defined with the **mc-method** command, is `local` or `xmlfile`.

Beyond an SSL proxy profile to secure communication to search the LDAP for the group name, use the following commands to complete the configuration.

- How to connect to the LDAP server. Use either of the following approaches.
 - The **mc-server-host** and **mc-server-port** commands
 - The **mc-loadbalancer-group** command
- Specify the login DN to access the LDAP server with the **mc-ldap-bind-dn** command
- Specify the user's password with the **mc-ldap-bind-password** command
- Optionally associate LDAP search parameters with the **mc-ldap-parameters** command

Example

Use a local XML file to map credentials and search the LDAP retrieve the distinguished name.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
# mc-ldap-search on
# mc-server-host ldap.mydomain.com
# mc-server-port 389
# mc-ldap-bind-dn "cn=proxyuser"
# mc-ldap-bind-password p@Ssw0rd
# mc-ldap-parameters ldap1-MC
```

Related reference

[“mc-ldap-bind-dn” on page 939](#)

This command specifies the login DN to access an LDAP server.

[“mc-ldap-bind-password” on page 933](#)

This command specifies the password for the login DN to access an LDAP server.

[“mc-ldap-search” on page 936](#)

This command indicates whether to retrieve the group names with an LDAP search.

[“mc-ldap-parameters” on page 934](#)

This command assigns the LDAP search parameters to search LDAP.

[“mc-loadbalancer-group” on page 938](#)

This command assigns a load balancer group as the target for an LDAP search.

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

[“mc-server-host” on page 941](#)

This command specifies the host of the LDAP server.

[“mc-server-port” on page 943](#)

This command specifies the listening port on the LDAP server.

mc-loadbalancer-group

This command assigns a load balancer group as the target for an LDAP search.

Syntax

mc-loadbalancer-group *name*

Parameters

name

Specifies the name of a load balancer group.

Guidelines

The **mc-loadbalancer-group** command assigns an LDAP load balancer group instead of a single LDAP server as the target for the search to retrieve the user's group.

This command is relevant only in the following situation:

- LDAP search is enabled with the **mc-ldap-search** command
- When the credentials mapping method, as defined with the **mc-method** command, is `local` or `xmlfile`.

This command is mutually exclusive with the combination of the **mc-server-host** and **mc-server-port** commands.

Beyond the LDAP load balancer group to search for the group name, use the following commands to complete the configuration.

- Optionally associate an SSL proxy profile with the **mc-ldap-sslproxy** command to secure communication with the LDAP server
- Specify the login DN to access the LDAP server with the **mc-ldap-bind-dn** command
- Specify the user's password with the **mc-ldap-bind-password** command
- Optionally associate LDAP search parameters with the **mc-ldap-parameters** command

Example

Uses a local XML file to map credentials and search LDAP to retrieve the DN.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
# mc-ldap-search on
# mc-loadbalancer-group LBGroup1
# mc-ldap-bind-dn "cn=proxyuser"
# mc-ldap-bind-password p@Ssw0rd
```

Related reference

[“mc-ldap-bind-dn” on page 939](#)

This command specifies the login DN to access an LDAP server.

[“mc-ldap-bind-password” on page 933](#)

This command specifies the password for the login DN to access an LDAP server.

[“mc-ldap-parameters” on page 934](#)

This command assigns the LDAP search parameters to search LDAP.

[“mc-ldap-sslproxy” on page 937](#)

This command assigns an SSL proxy profile with the LDAP credentials server.

[“mc-ldap-search” on page 936](#)

This command indicates whether to retrieve the group names with an LDAP search.

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

[“mc-server-host” on page 941](#)

This command specifies the host of the LDAP server.

[“mc-server-port” on page 943](#)

This command specifies the listening port on the LDAP server.

mc-ldap-bind-dn

This command specifies the login DN to access an LDAP server.

Syntax

mc-ldap-bind-dn *DN*

Parameters

DN

Specifies the login DN (distinguished name) to access the target LDAP server.

Guidelines

The **mc-ldap-bind-dn** command specifies the login DN to access the target LDAP server.

This command is relevant only in the following situation:

- LDAP search is enabled with the **mc-ldap-search** command
- When the credentials mapping method, as defined with the **mc-method** command, is `local` or `xmlfile`.

Beyond the login DN to search the LDAP for the group name, use the following commands to complete the configuration.

- How to connect to the LDAP server. Use either of the following approaches.
 - The **mc-server-host** and **mc-server-port** commands
 - The **mc-loadbalancer-group** command
- Optionally associate an SSL proxy profile with the **mc-ldap-sslproxy** command to secure communication with the LDAP server
- Specify the user's password with the **mc-ldap-bind-password** command

- Optionally associate LDAP search parameters with the **mc-ldap-parameters** command

Example

Use a local XML file to map credentials and search the LDAP retrieve the distinguished name.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
# mc-ldap-search on
# mc-server-host ldap.mydomain.com
# mc-server-port 389
# mc-ldap-bind-dn "cn=proxyuser"
# mc-ldap-bind-password p@Ssw0rd
# mc-ldap-parameters ldap1-MC
```

Related reference

[“mc-ldap-bind-password” on page 933](#)

This command specifies the password for the login DN to access an LDAP server.

[“mc-ldap-parameters” on page 934](#)

This command assigns the LDAP search parameters to search LDAP.

[“mc-ldap-search” on page 936](#)

This command indicates whether to retrieve the group names with an LDAP search.

[“mc-ldap-sslproxy” on page 937](#)

This command assigns an SSL proxy profile with the LDAP credentials server.

[“mc-loadbalancer-group” on page 938](#)

This command assigns a load balancer group as the target for an LDAP search.

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

[“mc-server-host” on page 941](#)

This command specifies the host of the LDAP server.

[“mc-server-port” on page 943](#)

This command specifies the listening port on the LDAP server.

mc-method

This command specifies the credentials mapping method.

Syntax

```
mc-method { custom | local | xmlfile }
```

Parameters

custom

Uses a custom style sheet. Use the **mc-custom-url** command to specify the location of the style sheet.

local

Uses the user group on the appliance. Does not access external resources. This setting is the default value and is the only option available if you have selected an **au-method** of local.

xmlfile

Uses an XML file on the appliance. Use the **mc-info-url** command to specify the location of the file.

Guidelines

The **mc-method** command sets the credential mapping (authorization) method for RBM.

The following table lists the supported credential mapping methods for each user authentication method.

Table 59. Authentication methods and supported credential mapping methods

Authentication method	Map credential method		
	local	xmlfile	custom
client-ssl	No	Yes	Yes
custom	No	Yes	Yes
ldap	No	Yes	Yes
local	Yes	Yes	Yes
radius	No	Yes	Yes
spnego	No	Yes	Yes
xmlfile	Yes	Yes	Yes

When the credentials mapping method is `local` or `xmlfile`, you can use the `mc-ldap-search` command to retrieve the distinguished name with an LDAP search.

Notes:

- The selected credentials mapping method must be fully configured before invoking this command.
- If the admin account is not configured with all permissions, the admin user is locked out of the GUI. Access the command line to change this circumstance.

Examples

- Set the authorization method to `xmlfile` and identifies the location of the file.

```
# mc-method xmlfile
# mc-info-url "local:///RBMPolicy.xml"
```

- Set the authorization method to `local`.

```
# mc-method local
```

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

[“mc-custom-url” on page 932](#)

This command specifies the URL of the RBM credential-mapping custom style sheet.

[“mc-info-url” on page 932](#)

This command specifies the URL of the mapping credentials XML file.

[“mc-ldap-search” on page 936](#)

This command indicates whether to retrieve the group names with an LDAP search.

mc-server-host

This command specifies the host of the LDAP server.

Syntax

mc-server-host *host*

Parameters

host

Specifies the IP address or domain name of the server.

Guidelines

The **mc-server-host** command specifies the IP address or domain name of the credentials server.

This command is relevant only in the following situation:

- LDAP search is enabled with the **mc-ldap-search** command
- When the credentials mapping method, as defined with the **mc-method** command, is `local` or `xmlfile`.

This command is mutually exclusive with the **mc-loadbalancer-group** command.

Beyond the LDAP server to search for the group name, use the following commands to complete the configuration.

- The **mc-server-port** command to specify the listening port on the LDAP server
- Optionally associate an SSL proxy profile with the **mc-ldap-sslproxy** command to secure communication with the LDAP server
- Specify the login DN to access the LDAP server with the **mc-ldap-bind-dn** command
- Specify the user's password with the **mc-ldap-bind-password** command
- Optionally associate LDAP search parameters with the **mc-ldap-parameters** command

Example

Use a local XML file to map credentials and search the LDAP retrieve the distinguished name.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
# mc-ldap-search on
# mc-server-host ldap.mydomain.com
# mc-server-port 389
# mc-ldap-bind-dn "cn=proxyuser"
# mc-ldap-bind-password p@Ssw0rd
# mc-ldap-parameters ldap1-MC
```

Related reference

[“mc-ldap-bind-dn” on page 939](#)

This command specifies the login DN to access an LDAP server.

[“mc-ldap-bind-password” on page 933](#)

This command specifies the password for the login DN to access an LDAP server.

[“mc-ldap-parameters” on page 934](#)

This command assigns the LDAP search parameters to search LDAP.

[“mc-ldap-sslproxy” on page 937](#)

This command assigns an SSL proxy profile with the LDAP credentials server.

[“mc-ldap-search” on page 936](#)

This command indicates whether to retrieve the group names with an LDAP search.

[“mc-loadbalancer-group” on page 938](#)

This command assigns a load balancer group as the target for an LDAP search.

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

[“mc-server-port” on page 943](#)

This command specifies the listening port on the LDAP server.

mc-server-port

This command specifies the listening port on the LDAP server.

Syntax

mc-server-port *port*

Parameters

port

Specifies the listening port on the server.

Guidelines

The **mc-server-port** command specifies the listening port on the LDAP server.

This command is relevant only in the following situation:

- LDAP search is enabled with the **mc-ldap-search** command
- When the credentials mapping method, as defined with the **mc-method** command, is `local` or `xmlfile`.

This command is mutually exclusive with the **mc-loadbalancer-group** command.

Beyond the listening port on the LDAP server to search the LDAP for the group name, use the following commands to complete the configuration.

- The **mc-server-host** command to specify the LDAP server
- Optionally associate an SSL proxy profile with the **mc-ldap-sslproxy** command to secure communication with the LDAP server
- Specify the login DN to access the LDAP server with the **mc-ldap-bind-dn** command
- Specify the user's password with the **mc-ldap-bind-password** command
- Optionally associate LDAP search parameters with the **mc-ldap-parameters** command

Example

Use a local XML file to map credentials and search the LDAP retrieve the distinguished name.

```
# mc-method xmlfile
# mc-info-url local:///RBM-MC.xml
# mc-ldap-search on
# mc-server-host ldap.mydomain.com
# mc-server-port 389
# mc-ldap-bind-dn "cn=proxyuser"
# mc-ldap-bind-password p@Ssw0rd
# mc-ldap-parameters ldap1-MC
```

Related reference

[“mc-ldap-bind-dn” on page 939](#)

This command specifies the login DN to access an LDAP server.

[“mc-ldap-bind-password” on page 933](#)

This command specifies the password for the login DN to access an LDAP server.

[“mc-ldap-parameters” on page 934](#)

This command assigns the LDAP search parameters to search LDAP.

[“mc-ldap-sslproxy” on page 937](#)

This command assigns an SSL proxy profile with the LDAP credentials server.

[“mc-ldap-search” on page 936](#)

This command indicates whether to retrieve the group names with an LDAP search.

[“mc-loadbalancer-group” on page 938](#)

This command assigns a load balancer group as the target for an LDAP search.

[“mc-method” on page 940](#)

This command specifies the credentials mapping method.

[“mc-server-host” on page 941](#)

This command specifies the host of the LDAP server.

password-hash-algorithm

This command sets the hash algorithm to apply to passwords before they are stored.

Syntax

password-hash-algorithm { [md5crypt](#) | sha256crypt }

Parameters

md5crypt

Uses MD5 Crypt as the hash algorithm. This setting is the default value.

sha256crypt

Uses SHA-256 Crypt as the hash algorithm.

Guidelines

The **password-hash-algorithm** command specifies the hash algorithm that is applied to passwords for locally defined users before the passwords are stored.

- In FIPS 140-2 Level 1 mode, the appliance cannot check MD5 Crypt password entries because MD5 is banned in this mode. If any existing account passwords use MD5 Crypt, the appliance refuses to enter FIPS 140-2 Level 1 mode to avoid user lockout. To successfully enter FIPS 140-2 Level 1 mode, you must select `sha256crypt` and then change the password on any existing user accounts that used MD5 Crypt when last changed.
- Firmware releases before 6.0.1 do not support SHA-256 Crypt passwords. If you need to downgrade to a release before 6.0.1, you must select `md5crypt` and then change the password on any existing user accounts that used SHA-256 Crypt when last changed. Only after such configuration is downgrading to the release before 6.0.1 allowed. This check is to avoid user lockout.

Example

Use the hash algorithm SHA-256 Crypt to apply to passwords before they are stored.

```
# password-hash-algorithm sha256crypt
```

pwd-aging

This command specifies whether users must periodically change their passwords.

Syntax

pwd-aging { on | [off](#) }

Parameters

on

Requires the periodic change of passwords.

off

Allows continued use of passwords. This setting is the default value.

Guidelines

If password-aging is enabled, use the **pwd-max-age** command to specify the maximum shelf-life of a user password.

Example

Require passwords to be changed every 15 days.

```
# pwd-aging on
# pwd-max-age 15
```

Related reference

[“pwd-max-age” on page 946](#)

This command specifies the maximum duration of passwords.

pwd-digit

This command specifies whether passwords must contain at least one numeric character.

Syntax

pwd-digit { on | off }

Parameters

on

Indicates that passwords must contain at least one numeric character.

off

Indicates that passwords do not require numeric characters. This setting is the default value.

Guidelines

When enabled, p4AssWord is acceptable, but password or PASSWORD is not acceptable.

When disabled, p4AssWord, password, or PASSWORD are acceptable.

pwd-history

This command specifies whether recent passwords can be reused.

Syntax

pwd-history { on | off }

Parameters

on

Indicates that passwords can be reused.

off

Indicates that passwords cannot be reused. This setting is the default value.

Guidelines

When enabled, use the **pwd-max-history** command to specify the number of passwords to retain. Passwords that are retained are not eligible for reuse.

Example

Indicate that the three most recent passwords cannot be reused.

```
# pwd-history on
# pwd-max-history 3
```

Related reference

[“pwd-max-history” on page 946](#)

This command specifies the number of passwords to retain.

pwd-max-age

This command specifies the maximum duration of passwords.

Syntax

pwd-max-age *days*

Parameters

days

Specifies the maximum number of days that a password is valid. Enter a value in the range 1 - 65535. The default value is 300.

Guidelines

If password-aging is enabled with the **pwd-aging** command, use the **pwd-max-age** command to specify the maximum shelf-life of a user password.

Example

Require passwords to be changed every 15 days.

```
# pwd-aging on
# pwd-max-age 15
```

Related reference

[“pwd-aging” on page 944](#)

This command specifies whether users must periodically change their passwords.

pwd-max-history

This command specifies the number of passwords to retain.

Syntax

pwd-max-history *count*

Parameters

count

Specifies the number of passwords to retain. Enter a value in the range 1 - 65535. The default value is 5.

Guidelines

If password reuse is enabled with the **pwd-history** command, use the **pwd-max-history** command to specify the number of recent passwords to retain. Passwords that are retained are not eligible for reuse.

Example

Indicate that the three most recent passwords cannot be reused.

```
# pwd-history on
# pwd-max-history 3
```

Related reference

[“pwd-history” on page 945](#)

This command specifies whether recent passwords can be reused.

pwd-minimum-length

This command specifies the minimum length of passwords.

Syntax

```
pwd-minimum-length length
```

Parameters

length

Specifies the minimum length. Enter a value in the range 1 - 128. The default value is 6.

pwd-mixed-case

This command specifies whether passwords must contain uppercase and lowercase characters.

Syntax

```
pwd-mixed-case { on | off }
```

Parameters

on

Indicates that passwords must contain uppercase and lowercase characters.

off

Indicates that passwords do not require uppercase and lowercase characters. This setting is the default value.

Guidelines

When enabled, pAssWord is acceptable, but password or PASSWORD is not acceptable.

When disabled, pAssWord, password, or PASSWORD is acceptable.

Examples

- Require passwords to contain both uppercase and lowercase characters.

```
# pwd-mixed-case on
```

- Restore the default state.

```
# pwd-mixed-case off
```

pwd-nonalphabetic

This command specifies whether passwords must contain nonalphabetic characters.

Syntax

```
pwd-nonalphabetic { on | off }
```

Parameters

on

Indicates that passwords must contain nonalphanumeric characters.

off

Indicates that passwords do not require nonalphanumeric characters. This setting is the default value.

Guidelines

When enabled, pa\$\$word is acceptable, but pAssWord or pa33word is not acceptable.

When disabled, pa\$\$word, pAssWord, or pa33word is acceptable.

Examples

- Require passwords to contain nonalphanumeric characters.

```
# pwd-nonalphanumeric on
```

- Restore the default state.

```
# pwd-nonalphanumeric off
```

pwd-username

This command specifies whether passwords can contain the user string.

Syntax

pwd-username { on | off }

Parameters

on

Indicates that passwords can contain the user name.

off

Indicates that passwords cannot contain the user name. This setting is the default value.

Guidelines

When enabled, the password BobPassword or password4Bob is acceptable for user Bob.

When disabled, the password BobPassword or password4Bob is not acceptable for user Bob.

Examples

- Allow passwords to contain the user.

```
# pwd-username on
```

- Restore the default state.

```
# pwd-username off
```

restrict-admin

This command specifies whether to restrict access by the `admin` account to the CLI through a serial connection.

Syntax

restrict-admin { on | off }

Parameters

on

Restricts the `admin` account to CLI access through a serial connection.

off

Allows the `admin` account to all access methods. This setting is the default value.

Guidelines

The **restrict-admin** command specifies whether to restrict access by the `admin` account to the CLI through a serial connection.

- When enabled, the access method for the `admin` account is through the CLI when connected through a serial connection.
- When disabled, the default state, the `admin` account can use all of the available access methods.

Note: On XI50z, this option is disabled. No administrator can modify this setting.

Examples

- Restrict CLI access by the `admin` account to serial connections.

```
# restrict-admin on
```

- Allow access by the `admin` account to all access methods.

```
# restrict-admin off
```

ssl-client

This command associates an SSL client profile for LDAP authentication.

Syntax

ssl-client *name*

no ssl-client

Parameters

name

Specifies the name of an SSL client profile.

Guidelines

The **ssl-client** command specifies the SSL client profile to secure connections with the LDAP server during LDAP authentication.

To create an SSL client profile, use the **Crypto ssl-client** command. To remove the SSL client profile, use the **no ssl-client** command.

This command is relevant when the following conditions are met.

- The authentication method set by the **au-method** command is ldap.
- The type set by the **ssl-client-type** command is client.

Related reference

[“au-method” on page 923](#)

This command specifies the RBM authentication method.

ssl-client-type

This command sets the type of the SSL profile for LDAP authentication.

Syntax

```
ssl-client-type { proxy | client }
```

Parameters

proxy

This value is deprecated. Do not use.

client

Uses the SSL client profile to secure connections.

Guidelines

The **ssl-client-type** command sets the SSL profile type to secure connections with the LDAP server during LDAP authentication. To specify an SSL client profile, use the **ssl-client** command.

Related reference

[“ssl-client” on page 949](#)

This command associates an SSL client profile for LDAP authentication.

REST Management Interface commands

Use the REST Management Interface commands to enable and modify the configuration of the REST management interface.

You enter the commands in REST Management Interface configuration mode:

1. From the appliance command line, enter global configuration mode:

```
config
```

2. From global configuration mode, type `rest-mgmt` to enter REST Management Interface configuration mode.
3. Type `exit` to leave the configuration mode and save your changes, then type `exit` again to leave global configuration mode.

local-address

This command assigns a local IP address on which the REST management interface listens.

Syntax

```
local-address address
```

Parameters

address

Identifies the IP address on which the appliance listens for incoming REST management requests.

Guidelines

The **local-address** command assigns a local IP address on which the REST management interface listens.

You can specify the local address and port together with the **local-address** command or specify the port independently with the **port** command.

Examples

Specify a listening address for the REST management interface:

```
# local-address 192.0.2.2
```

Specify a listening address and port for the REST management interface:

```
# local-address 192.0.2.2 5552
```

Related reference

[“port” on page 951](#)

This command assigns the local port on which the REST management interface listens.

port

This command assigns the local port on which the REST management interface listens.

Syntax

port *port*

Parameters

port

Identifies the port on the appliance. The default value is 5554.

Guidelines

The **port** command assigns the local port on which the REST management interface listens.

To specify the local IP address for the REST management interface, use the **local-address** command.

Examples

Specify a listening port for the REST management interface.

```
# port 5552
```

Related reference

[“local-address” on page 950](#)

This command assigns a local IP address on which the REST management interface listens.

ssl-config-type

This command sets the type of the SSL profile for the REST management interface.

Syntax

ssl-config-type *config_type*

Parameters

config_type

Specify `server` if you want to specify SSL (TLS) security for the REST management interface.

Guidelines

The **ssl-config-type** command sets the SSL profile type to secure connections between clients and the appliance. You must use the `server` type.

Related reference

[“ssl-server” on page 952](#)

This command associates an SSL server profile with the REST management interface.

ssl-server

This command associates an SSL server profile with the REST management interface.

Syntax

ssl-server *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`.

SNMP Settings commands

SNMP Settings mode provides the commands to modify the SNMP settings.

To enter the mode, use the Global **snmp** command. To disable SNMP, use the Global **no snmp** command.

While in this mode, use the following commands to define the server connection. The settings grant access to the SNMP agent for an SNMP manager, identify the appliance UDP port that is monitored by the SNMP agent, and manages the trap events from the SNMP agent.

- 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.

Related reference

[“snmp” on page 816](#)

This command enters SNMP Settings mode.

access-level

This command specifies the level of access that an SNMPv3 manager has to the appliance MIBs.

Syntax

access-level `read-only`

access-level `read-write`

Parameters

read-only

Indicates that managers are restricted to SNMP **get** operations, which means that these managers can read, but cannot change management information base (MIB) values.

read-write

Indicates that managers have access to both SNMP **get** and **set** operations, which means that these managers can read and change MIB values.

Examples

- Specify read-only access.

```
# access-level read-only
```

community

This command grants and defines access to the specified SNMPv1 or v2c communities.

Syntax

```
community communityName read-only [IP_address]
```

```
community communityName read-write [IP_address]
```

Parameters

communityName

Specifies the name of the community. The community name is effectively a password phrase that accompanies SNMP requests, and is used to determine whether the request can be fulfilled or not.

read-only

Indicates that managers are restricted to SNMP **get** operations, which means that these managers can read, but cannot change management information base (MIB) values.

read-write

Indicates that managers have access to both SNMP **get** and **set** operations, which means that these managers can read and change MIB values.

IP_address

Specify an IP address to restrict access to the SNMP manager in the named community with the specified IP address. By default, any SNMP manager belonging to the named community can make requests.

Examples

- Specify read-only access to the community named "private".

```
# community private read-only
```

ip-address

This command specifies the IP address that is listened on for SNMP requests.

Syntax

```
ip-address local_IP_address
```

Parameters

local_IP_address

Specify a local IP address that the SNMP service listens on for SNMP requests to the appliance. Specify 0.0.0.0 to listen on all appliance interfaces.

Examples

- Specify the appliance listens on IP address 198.51.100.0 for SNMP requests.

```
# ip-address 198.51.100.0
```

port

This command identifies the appliance UDP port that is monitored by the SNMP agent for SNMP requests.

Syntax

port [*address*] *port*

Parameters

address

Specifies an optional IP address that identifies a specific local interface as a recipient of SNMP requests.

port

Identifies the UDP port that is monitored by the SNMP agent or engine for SNMP requests. The default value is 161.

Guidelines

In the absence of an IP address argument, the SNMP agent monitors the specified port on all interfaces for SNMP requests. With an IP address argument provided, the SNMP agent monitors only the specified interface-port pair for SNMP requests. For example, the IP address might be the XML management port.

Examples

- Identify a nonstandard SNMP port, 65161, on all interfaces as the recipient of SNMP requests.

```
# port 65161
```

- Identify a specific IP address and port pair as the recipient of SNMP requests.

```
# port 10.10.10.10 161
```

security-level

This command specifies the access that an SNMPv3 manager has to the appliance.

Syntax

security-level noAuthNoPriv

security-level authNoPriv

security-level authPriv

Parameters

noAuthNoPriv

The SNMP connection requires neither authentication of users nor encryption of data.

authNoPriv

The SNMP connection requires authentication of users but not the encryption of data.

authPriv

The SNMP connection requires authentication of users and encryption of data.

Examples

- Specify that the appliance requires an SNMPv3 manager to supply valid user and credentials, and that data transferred to the manager is encrypted.

```
# security-level authPriv
```

trap-code

This command adds an event code to the trap list.

Syntax

trap-code *code*

no trap-code *code*

Parameters

code

Specifies the hex identifier of an event code.

Guidelines

The **trap-code** command specifies individual event codes to add to the trap list. Run this command for each event to add to the list.

Use the **no trap-code** command to delete a previously configured code from the trap list.

The following is a list of the event subscriptions that cannot be deleted. Even if you delete one or more these subscriptions, they are added after the next restart.

- 0x00030002 (Out of memory)
- 0x00230003 (Unable to allocate execution resources)
- 0x00330002 (Memory full)
- 0x00b30014 (Duplicate IP address)
- 0x00e30001 (NTP - Cannot Resolve Server Name)
- 0x00e40008 (NTP Timeout Error)
- 0x00f30008 (File is expired)
- 0x01530001 (Time zone config mismatch.)
- 0x01a2000e (Installed battery is nearing end of life.)
- 0x01a40001 (Throttling connections due to low memory)
- 0x01a40005 (Throttling connections due to low temporary file space)
- 0x01a40008 (Throttling connections due to low number of free ports)
- 0x01b10009 (uncertified HSM firmware detected)
- 0x01b20002 (HSM is uninitialized)
- 0x01b20004 (HSM PED login failed)
- 0x01b20008 (HSM password login failed)
- 0x02220001 (Power supply failure.)
- 0x02220003 (Internal cooling fan has stopped.)
- 0x02240002 (Internal cooling fan has slowed)

Note: The "File is expired" event refers to the file associated with the certificate aliases on the appliance.

This command is relevant when the **trap-default-subscriptions** property is enabled.

Examples

Add the "Out of memory" parser event with hex code 0x00030002 to the list.

```
# trap-code 0x00030002
```

Related reference

[“trap-default-subscriptions” on page 956](#)

This command enables or disables the default list of event codes that generate traps.

trap-default-subscriptions

This command enables or disables the default list of event codes that generate traps.

Syntax

trap-default-subscriptions { on | off }

Parameters

on

The default event subscriptions are used. This setting is the default value.

off

The default event subscriptions are not used. Your event subscriptions persist after a system restart.

Guidelines

The **trap-default-subscriptions** command controls whether you can define your own event subscriptions.

- When enabled, use the **trap-code** and **trap-priority** command to define subscriptions.
- When disabled, define an SNMP log target.

Related reference

[“trap-code” on page 955](#)

This command adds an event code to the trap list.

[“trap-priority” on page 956](#)

This command specifies the minimum criticality for trap events.

trap-priority

This command specifies the minimum criticality for trap events.

Syntax

trap-priority *priority*

Parameters

priority

Identifies the event priority. The default value is **error**.

Guidelines

The **trap-priority** command specifies the minimum criticality for trap events. The priorities are hierarchical. In descending order of criticality, the priorities are **emergency**, **alert**, **critic**, **error**, **warn**, **notice**, **info**, and **debug**.

This command is relevant when the **trap-default-subscriptions** command is set to on.

Example

Set the trap priority to warn or greater criticality.

```
# trap-priority warn
```

Related reference

[“trap-default-subscriptions” on page 956](#)

This command enables or disables the default list of event codes that generate traps.

trap-target

This command specifies the recipient of SNMP traps issued by the local SNMP agent.

Syntax

For SNMP v3:

trap-target *IPaddress port 3 user security_level*

For SNMP v1:

trap-target *IPaddress port 1 community*

For SNMP v2c:

trap-target *IPaddress port 2c community*

Parameters

IPaddress

Specifies the IP address that receives traps/notifications.

port

Identifies a UDP port at the IP address. The default value is 162.

user

The user ID to authenticate with when sending notifications in SNMP v3.

security_level

The SNMP security level, is one of:

- noAuthNoPriv (no user authentication, no data encryption)
- authNoPriv (user authentication, no data encryption)
- authPriv (user authentication, data encryption)

community

Provides a community name, which is essentially a password, to include in the SNMP message header. The default value is public.

Guidelines

The local SNMP agent or engine issues the following generic traps:

- coldStart
- linkDown
- linkUp
- authenticationFailure

Use the **no trap-target** command to delete a previously configured recipient of SNMP traps.

Examples

- Specify a recipient of SNMP v3 traps at 10.10.10.20:162. The trap recipient is accessed with user ID "snmpNs", and authentication and encryption are required.

```
# trap-target 10.10.10.20 3 snmpNs authPriv
```

- Specify a recipient of SNMP v2c traps at 10.10.10.11:162. The trap recipient is accessed with the public community.

```
# trap-target 10.10.10.11 2c
```

- Specify a recipient of SNMP traps at 10.10.100.19:162. The trap recipient is accessed with the OpenView community.

```
# trap-target 10.10.100.19 2c OpenView
```

- Delete 10.10.10.11:162 from the list of trap recipients.

```
# no trap-target 10.10.10.11
```

user

This command specifies the user ID used by a SNMPv3 manager to connect to the appliance.

Syntax

user *userName*

Parameters

userName

The name of a local user on the appliance. The local user must have been defined with SNMPv3 credentials, see [“snmp-cred” on page 721](#).

version

This command specifies the supported SNMP version.

Syntax

version { 1 | 2c | 3 }

Parameters

1

Specifies support for SNMP Version 1 as defined in RFC 1155, RFC 1156, and RFC 1157.

2c

Specifies support for SNMP Version 2c as originally defined in RFC 1901 through RFC 1908. This setting is the default value.

3

Specifies support for SNMP Version 3 as defined in RFC 2261 through RFC 2265.

Examples

- Specifies support for SNMP Version 1.

```
# version 1
```

- Specifies support for SNMP Version 2c, which is, the default state.

```
# version 2c
```

SSH Server Profile commands

SSH Server Profile mode provides the commands to manage the cipher suites for SSH connections.

To enter the mode, use the Crypto **sshserverprofile** command.

ciphers

This command specifies the cipher suites that the appliance accepts for SSH encryption.

Syntax

Add a cipher.

ciphers *cipher-string*

Delete a cipher.

no ciphers *cipher-string*

Parameters

cipher-string

Specifies the ciphers allowed by OpenSSH version 2 to use in SSH communication. The order of cipher suites is important. The server compares its list to the client's list in order of preference. The first cipher suite in the client's list is chosen when it is also supported by the server.

The cipher suites in the following list are supported and all are default values. You can change the preference order of cipher suites.

- CHACHA20-POLY1305_AT_OPENSSSH.COM
- AES128-CTR
- AES192-CTR
- AES256-CTR
- AES128-GCM_AT_OPENSSSH.COM
- AES256-GCM_AT_OPENSSSH.COM
- ARCFOUR256
- ARCFOUR128
- AES128-CBC
- 3DES-CBC
- BLOWFISH-CBC
- CAST128-CBC
- AES192-CBC
- AES256-CBC
- ARCFOUR
- RIJNDAEL-CBC_AT_LYSATOR.LIU.SE

SSL Client Profile commands

SSL Client Profile mode provides the commands to create or modify an SSL client profile.

To enter the mode, use the **Crypto ssl-client** command. To delete an SSL client profile, use the **Crypto no ssl-client** command.

While in this mode, use the following commands to define the SSL client profile.

- 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.

cache-size

This command specifies the maximum number of client sessions to cache.

Syntax

cache-size *entries*

Parameters

entries

Specifies the maximum number of client sessions to cache. Enter a value in the range 1 - 500000. The default value is 100.

Guidelines

The cache-size command specifies the maximum number of client sessions to cache.

This command is relevant only when the **catching** command is set to on.

cache-timeout

This command sets the time that SSL sessions remain in the client session cache before they are removed.

Syntax

cache-timeout *seconds*

Parameters

seconds

Sets the time that SSL sessions remain in the SSL session cache. Enter a value in the range 1 - 86400. The default value is 300.

Guidelines

The cache-timeout command sets the time that SSL sessions remain in the session cache before they are removed.

This command is relevant only when the **catching** command is set to on.

catching

This command controls whether to cache SSL sessions when the appliance is an SSL client.

Syntax

catching { on | off }

Parameters

on

Enables SSL session caching. This setting is the default value.

off

Disables SSL session caching.

Guidelines

The caching command controls whether to cache SSL sessions. When enabled, you can change the following SSL session settings. The default behavior is 100 entries and 5 minutes.

- Use the **cache-size** command to specify the maximum number of SSL sessions to cache.
- Use the **cache-timeout** command to specify the time that SSL sessions remain in the SSL session cache before they are removed.

ciphers

This command specifies the cipher suites that the SSL client profile uses to establish a secure connection.

Syntax

ciphers *cipher_string*

Parameters

cipher_string

Specifies the cipher suites. The following cipher suites are supported.

- RSA_WITH_NULL_MD5
- RSA_WITH_NULL_SHA
- RSA_EXPORT_WITH_RC4_40_MD5
- RSA_WITH_RC4_128_MD5
- RSA_WITH_RC4_128_SHA
- RSA_EXPORT_WITH_RC2_CBC_40_MD5
- RSA_EXPORT_WITH_DES40_CBC_SHA
- RSA_WITH_DES_CBC_SHA
- RSA_WITH_3DES_EDE_CBC_SHA (default)
- DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
- DHE_DSS_WITH_DES_CBC_SHA
- DHE_DSS_WITH_3DES_EDE_CBC_SHA (default)
- DHE_RSA_EXPORT_WITH_DES40_CBC_SHA (default)
- DHE_RSA_WITH_DES_CBC_SHA
- DHE_RSA_WITH_3DES_EDE_CBC_SHA (default)
- RSA_WITH_AES_128_CBC_SHA (default)
- DHE_DSS_WITH_AES_128_CBC_SHA (default)
- DHE_RSA_WITH_AES_128_CBC_SHA (default)
- RSA_WITH_AES_256_CBC_SHA (default)
- DHE_DSS_WITH_AES_256_CBC_SHA (default)
- DHE_RSA_WITH_AES_256_CBC_SHA (default)
- RSA_WITH_NULL_SHA256
- RSA_WITH_AES_128_CBC_SHA256 (default)
- RSA_WITH_AES_256_CBC_SHA256 (default)
- DHE_DSS_WITH_AES_128_CBC_SHA256 (default)
- DHE_RSA_WITH_AES_128_CBC_SHA256 (default)
- DHE_DSS_WITH_AES_256_CBC_SHA256 (default)
- DHE_RSA_WITH_AES_256_CBC_SHA256 (default)

- RSA_WITH_AES_128_GCM_SHA256 (default)
- RSA_WITH_AES_256_GCM_SHA384 (default)
- DHE_RSA_WITH_AES_128_GCM_SHA256 (default)
- DHE_RSA_WITH_AES_256_GCM_SHA384 (default)
- DHE_DSS_WITH_AES_128_GCM_SHA256 (default)
- DHE_DSS_WITH_AES_256_GCM_SHA384 (default)ECDHE_RSA_WITH_NULL_SHA
- ECDHE_RSA_WITH_RC4_128_SHA
- ECDHE_RSA_WITH_3DES_EDE_CBC_SHA (default)
- ECDHE_RSA_WITH_AES_128_CBC_SHA
- ECDHE_RSA_WITH_AES_256_CBC_SHA
- ECDHE_RSA_WITH_AES_128_CBC_SHA256 (default)
- ECDHE_RSA_WITH_AES_256_CBC_SHA384 (default)
- ECDHE_RSA_WITH_AES_128_GCM_SHA256 (default)
- ECDHE_RSA_WITH_AES_256_GCM_SHA384 (default)
- ECDHE_ECDSA_WITH_NULL_SHA
- ECDHE_ECDSA_WITH_RC4_128_SHA
- ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA (default)
- ECDHE_ECDSA_WITH_AES_128_CBC_SHA (default)
- ECDHE_ECDSA_WITH_AES_256_CBC_SHA (default)
- ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 (default)
- ECDHE_ECDSA_WITH_AES_256_CBC_SHA384 (default)
- ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 (default)
- ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 (default)

Guidelines

The ciphers command specifies the cipher suites that the SSL client profile uses to establish a secure connection.

The cipher suites correspond to the RFC names without the TLS_ or SSL_ prefix. For example, RSA_WITH_3DES_EDE_CBC_SHA correspond to TLS_RSA_WITH_3DES_EDE_CBC_SHA or SSL_RSA_WITH_3DES_EDE_CBC_SHA in the relevant RFC.

The SSL client profile must include at least one cipher suite that matches the associated key material.

- An RSA signing key requires ECDHE_RSA cipher suites.
- An ECDSA signing key requires ECDHE_ECDSA cipher suites.

To specify multiple cipher suites, run this command for each cipher suite.

curves

This command specifies the list of elliptic curves that the SSL client profile supports.

Syntax

curves *name*

Parameters

name

Specifies the name of the curve. The following curves are supported:

- sect163k1
- sect163r1
- sect163r2
- sect193r1
- sect193r2
- sect233k1
- sect233r1
- sect239k1
- sect283k1
- sect283r1
- sect409k1
- sect409r1
- sect571k1
- sect571r1
- secp160k1
- secp160r1
- secp160r2
- secp192k1
- secp192r1
- secp224k1
- secp224r1
- secp256k1 (default)
- secp256r1 (default)
- secp384r1 (default)
- secp521r1 (default)
- brainpoolP256r1 (default)
- brainpoolP384r1 (default)
- brainpoolP512r1 (default)

Guidelines

The **curves** command specifies the elliptic curves that the SSL client profile supports.

To specify multiple curves, run this command for each curve.

idcred

This command specifies the identification credentials the appliance uses to authenticate itself to an SSL server if the SSL server requests client authentication.

Syntax

idcred *name*

Parameters

name

Specifies the name of the identification credentials.

Guidelines

The **idcred** command specifies the identification credentials that the appliance uses to authenticate itself to an SSL server if client authentication is requested.

To create identification credentials, use the **Crypto idcred** command.

protocols

This command specifies the SSL and TLS protocol versions to support when the appliance is an SSL client.

Syntax

protocols *option-string*

Parameters

option-string

Specifies the SSL and TLS protocol versions to support. When you enable support for multiple protocol versions, use a plus sign (+) character to separate the versions. The following values are valid. The default value is TLSv1d0+TLSv1d1+TLSv1d2.

SSLv3

Enables SSL 3.0.

TLSv1d0

Enables TLS 1.0.

TLSv1d1

Enables TLS 1.1.

TLSv1d2

Enables TLS 1.2.

Guidelines

The **protocols** command specifies the SSL and TLS protocol versions to support.

ssl-client-features

This command specifies the features to add in the SSL client profile.

Syntax

ssl-client-features *feature*

Parameters

feature

Identifies the feature. Separate multiple features with the plus sign (+) character. The default value is use-sni.

use-sni

Allows the client to send the Server Name Indication (SNI) extension in the ClientHello message to the server that the client attempts to connect to.

permit-insecure-servers

Allows connections to SSL servers that do not support RFC 5746.

compression

Enables SSL compression. Compression in HTTPS is dangerous because the HTTPS connection becomes vulnerable to the CRIME (Compression Ratio Info-leak Made Easy) attack.

Guidelines

The **ssl-client-features** command specifies the features to add in the SSL client profile that secures connections between the appliance and its targets.

valcred

This command specifies the name of the validation credentials to validate the SSL server certificate during the SSL handshake.

Syntax

valcred *valcred*

Parameters

valcred

Specifies the name of the validation credentials.

Guidelines

The **valcred** command specifies the name of the validation credentials to validate the SSL server certificate during the SSL handshake. Validation credentials consist of a set of certificates that validates the authenticity of received certificates and digital signatures.

This command is required when the `validate-server-cert` command is set to on.

validate-server-cert

This command controls whether to validate the server certificate during the SSL handshake.

Syntax

validate-server-cert { on | off }

Parameters

on

Validates the server certificate. This setting is the default value.

off

Does not validate the server certificate.

Guidelines

The **validate-server-cert** command controls whether to validate the server certificate during the SSL handshake. When enabled, use the **valcred** command to specify the validation credentials with the certificate to validate the server certificate.

SSL Server Profile commands

You can use the SSL Server Profile commands to create or modify an SSL server profile.

To enter SSL Server Profile mode, use the **crypto ssl-server** command. To delete an SSL server profile, use the **crypto no ssl-server** command.

While in this mode, use the commands in the following table to define the SSL server profile.

- 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.

allow-legacy-renegotiation

This command controls whether to allow SSL renegotiation with SSL clients that do not support RFC 5746.

Syntax

allow-legacy-renegotiation { on | off }

Parameters

on

Allows SSL renegotiation with SSL clients that do not support RFC 5746.

off

Does not allow SSL renegotiation with SSL clients that do not support RFC 5746. This setting is the default value.

Guidelines

The **allow-legacy-renegotiation** command controls whether to allow SSL renegotiation with SSL clients that do not support RFC 5746. By default, this support is disabled because renegotiation with such clients is vulnerable to man-in-the-middle attacks as documented in CVE-2009-3555.

cache-size

This command specifies the maximum number of server sessions to cache.

Syntax

cache-size *entries*

Parameters

entries

Specifies the maximum number of server sessions to cache in kilo entries (1024 entries). Enter a value in the range 1 - 500. The default value is 20.

Guidelines

The cache-size command specifies the maximum number of server sessions to cache.

This command is relevant only when the caching command is set to on.

Related reference

[“caching” on page 967](#)

This command controls whether to cache the SSL sessions when the appliance is an SSL server.

cache-timeout

This command sets the time that SSL sessions remain in the server session cache before they are removed.

Syntax

cache-timeout *seconds*

Parameters

seconds

Sets the time that SSL sessions remain in the SSL session cache. Enter a value in the range 1 - 86400. The default value is 300.

Guidelines

The **cache-timeout** command sets the time that SSL sessions remain in the session cache before they are removed.

Related reference

[“caching” on page 967](#)

This command controls whether to cache the SSL sessions when the appliance is an SSL server.

caching

This command controls whether to cache the SSL sessions when the appliance is an SSL server.

Syntax

caching { on | off }

Parameters

on

Enables SSL session caching. This setting is the default value.

off

Disables SSL session caching.

Guidelines

The caching command controls whether to cache SSL sessions. When enabled, you can change the following SSL session settings. The default behavior is 20480 entries and 5 minutes.

- Use the **cache-size** command to specify the maximum number of SSL sessions to cache.
- Use the **cache-timeout** command to specify the time that SSL sessions remain in the SSL session cache before they are removed.

Related reference

[“cache-size” on page 966](#)

This command specifies the maximum number of server sessions to cache.

[“cache-timeout” on page 966](#)

This command sets the time that SSL sessions remain in the server session cache before they are removed.

ciphers

This command specifies the cipher suites that the SSL server profile uses to establish a secure connection.

Syntax

ciphers *cipher_string*

Parameters

cipher_string

Specifies the cipher suites. The following cipher suites are supported.

- RSA_WITH_NULL_MD5

- RSA_WITH_NULL_SHA
- RSA_EXPORT_WITH_RC4_40_MD5
- RSA_WITH_RC4_128_MD5
- RSA_WITH_RC4_128_SHA
- RSA_EXPORT_WITH_RC2_CBC_40_MD5
- RSA_EXPORT_WITH_DES40_CBC_SHA
- RSA_WITH_DES_CBC_SHA
- RSA_WITH_3DES_EDE_CBC_SHA (default)
- DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
- DHE_DSS_WITH_DES_CBC_SHA
- DHE_DSS_WITH_3DES_EDE_CBC_SHA (default)
- DHE_RSA_EXPORT_WITH_DES40_CBC_SHA (default)
- DHE_RSA_WITH_DES_CBC_SHA
- DHE_RSA_WITH_3DES_EDE_CBC_SHA (default)
- RSA_WITH_AES_128_CBC_SHA (default)
- DHE_DSS_WITH_AES_128_CBC_SHA (default)
- DHE_RSA_WITH_AES_128_CBC_SHA (default)
- RSA_WITH_AES_256_CBC_SHA (default)
- DHE_DSS_WITH_AES_256_CBC_SHA (default)
- DHE_RSA_WITH_AES_256_CBC_SHA (default)
- RSA_WITH_NULL_SHA256
- RSA_WITH_AES_128_CBC_SHA256 (default)
- RSA_WITH_AES_256_CBC_SHA256 (default)
- DHE_DSS_WITH_AES_128_CBC_SHA256 (default)
- DHE_RSA_WITH_AES_128_CBC_SHA256 (default)
- DHE_DSS_WITH_AES_256_CBC_SHA256 (default)
- DHE_RSA_WITH_AES_256_CBC_SHA256 (default)
- RSA_WITH_AES_128_GCM_SHA256 (default)
- RSA_WITH_AES_256_GCM_SHA384 (default)
- DHE_RSA_WITH_AES_128_GCM_SHA256 (default)
- DHE_RSA_WITH_AES_256_GCM_SHA384 (default)
- DHE_DSS_WITH_AES_128_GCM_SHA256 (default)
- DHE_DSS_WITH_AES_256_GCM_SHA384 (default)ECDHE_RSA_WITH_NULL_SHA
- ECDHE_RSA_WITH_RC4_128_SHA
- ECDHE_RSA_WITH_3DES_EDE_CBC_SHA (default)
- ECDHE_RSA_WITH_AES_128_CBC_SHA
- ECDHE_RSA_WITH_AES_256_CBC_SHA
- ECDHE_RSA_WITH_AES_128_CBC_SHA256 (default)
- ECDHE_RSA_WITH_AES_256_CBC_SHA384 (default)
- ECDHE_RSA_WITH_AES_128_GCM_SHA256 (default)
- ECDHE_RSA_WITH_AES_256_GCM_SHA384 (default)
- ECDHE_ECDSA_WITH_NULL_SHA
- ECDHE_ECDSA_WITH_RC4_128_SHA

- ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA (default)
- ECDHE_ECDSA_WITH_AES_128_CBC_SHA (default)
- ECDHE_ECDSA_WITH_AES_256_CBC_SHA (default)
- ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 (default)
- ECDHE_ECDSA_WITH_AES_256_CBC_SHA384 (default)
- ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 (default)
- ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 (default)

Guidelines

The `ciphers` command specifies the cipher suites that the SSL server profile uses to establish a secure connection.

The cipher suites correspond to the RFC names without the `TLS_` or `SSL_` prefix. For example, `RSA_WITH_3DES_EDE_CBC_SHA` correspond to `TLS_RSA_WITH_3DES_EDE_CBC_SHA` or `SSL_RSA_WITH_3DES_EDE_CBC_SHA` in the relevant RFC.

The SSL server profile must include at least one cipher suite that matches the associated key material.

- An RSA signing key requires ECDHE_RSA cipher suites.
- An ECDSA signing key requires ECDHE_ECDSA cipher suites.

The SSL server profile must include at least one cipher suite that matches the identification credentials as specified by the `idcred` command.

- When the identification credentials contains RSA keys, you must specify at least one RSA cipher suite.
- When the identification credentials contains ECDSA keys, you must specify at least one ECDSA cipher suite.

To specify multiple cipher suites, run this command for each cipher suite.

Related reference

[“idcred” on page 971](#)

This command specifies the identification credentials that authenticate the appliance during the SSL handshake.

compression

This command controls whether to enable SSL compression when the appliance is an SSL server.

Syntax

compression { on | off }

Parameters

on

Enables SSL compression

off

Disables SSL compression. This setting is the default value.

Guidelines

The `compression` command controls whether to enable SSL compression. Compression in HTTPS is dangerous because the HTTPS connection becomes vulnerable to the CRIME (Compression Ratio Info-leak Made Easy) attack.

curves

This command specifies the list of elliptic curves that the SSL server profile supports.

Syntax

curves *name*

Parameters

name

Specifies the name of the curve. The following curves are supported:

- sect163k1
- sect163r1
- sect163r2
- sect193r1
- sect193r2
- sect233k1
- sect233r1
- sect239k1
- sect283k1
- sect283r1
- sect409k1
- sect409r1
- sect571k1
- sect571r1
- secp160k1
- secp160r1
- secp160r2
- secp192k1
- secp192r1
- secp224k1
- secp224r1
- secp256k1 (default)
- secp256r1 (default)
- secp384r1 (default)
- secp521r1 (default)
- brainpoolP256r1 (default)
- brainpoolP384r1 (default)
- brainpoolP512r1 (default)

Guidelines

The **curves** command specifies the elliptic curves that the SSL server profile supports.

To specify multiple curves, run this command for each curve.

idcred

This command specifies the identification credentials that authenticate the appliance during the SSL handshake.

Syntax

idcred *name*

Parameters

name

Specifies the name of the identification credentials.

Guidelines

The **idcred** command specifies the identification credentials that authenticate the appliance during the SSL handshake.

To create identification credentials, use the **Crypto idcred** command.

Related reference

[“idcred” on page 733](#)

This command creates identification credentials.

max-duration

This command specifies the maximum time to maintain an SSL session after the initial negotiated handshake when the appliance is an SSL server.

Syntax

max-duration *seconds*

Parameters

seconds

Specifies the maximum time in seconds. Enter a value in the range 1 - 11520. The default value is 60.

Guidelines

The **max-duration** command specifies the maximum time to maintain an SSL session after the initial negotiated handshake.

This command is relevant only when the value set by the **ssl-options** command contains **max-duration**.

Related reference

[“ssl-options” on page 975](#)

This command specifies the list of elliptic curves that the SSL server profile supports.

max-renegotiation-allowed

This command specifies the maximum number of renegotiation attempts that a client can initiate per session.

Syntax

max-renegotiation-allowed *seconds*

Parameters

seconds

Specifies the maximum number of renegotiation attempts that a client can initiate per session. Enter a value in the range 0 - 512. The default value is 0, which disables the support for the client initiated renegotiation.

Guidelines

The **max-renegotiation-allowed** command specifies the maximum number of renegotiation attempts that a client can initiate per session. The client can initiate renegotiation after the original negotiation in the handshake.

This command is relevant only when the value set by the **ssl-options** command contains max-renegotiation.

Related reference

[“ssl-options” on page 975](#)

This command specifies the list of elliptic curves that the SSL server profile supports.

prefer-server-ciphers

This command controls whether to use the server's cipher suite order instead of the client's during cipher suite negotiation.

Syntax

prefer-server-ciphers { on | off }

Parameters

on

Uses the server's cipher suite order. This setting is the default value.

off

Uses the client's cipher suite order.

Guidelines

The prefer-server-ciphers command controls whether to use the server's cipher suite order instead of the client's cipher suite order during cipher suite negotiation.

When the server and the client negotiate which cipher suites to use, the server compares the client's cipher suites list with the server's cipher suites list and select a shared one for connection. For example:

- The client's cipher suite list in order of the client's preference:
ECDHE_RSA_WITH_3DES_EDE_CBC_SHA, DHE_RSA_WITH_AES_256_CBC_SHA256,
RSA_WITH_AES_256_CBC_SHA256
- The server's cipher suite list in order of the server's preference:
DHE_RSA_WITH_AES_256_CBC_SHA256, RSA_WITH_DES_CBC_SHA,
ECDHE_RSA_WITH_3DES_EDE_CBC_SHA

The following rules apply.

- When enabled, the server selects the first shared cipher suite based on the server's preference. In this case, DHE_RSA_WITH_AES_256_CBC_SHA256 is used.
- When disabled, the server selects the first shared cipher suite based on the client's preference. In this case, ECDHE_RSA_WITH_3DES_EDE_CBC_SHA is used.

prohibit-resume-on-reneg

This command controls whether a previous SSL session can be resumed during a renegotiation handshake.

Syntax

prohibit-resume-on-reneg { on | off }

Parameters

on

Indicates that a previous session cannot be resumed during a renegotiation handshake.

off

Indicates that a previous session can be resumed during a renegotiation handshake. This setting is the default value.

Guidelines

The **prohibit-resume-on-reneg** command controls whether a previous SSL session can be resumed during a renegotiation handshake.

This command is relevant only when the value set by the **max-renegotiation-allowed** command is not 0.

Related reference

[“max-renegotiation-allowed” on page 971](#)

This command specifies the maximum number of renegotiation attempts that a client can initiate per session.

protocols

This command specifies the SSL and TLS protocol versions to support when the appliance is an SSL server.

Syntax

protocols *option-string*

Parameters

option-string

Specifies the SSL and TLS protocol versions to support. When you enable support for multiple protocol versions, use a plus sign (+) character to separate the versions. The following values are valid. The default value is TLSv1d0+TLSv1d1+TLSv1d2.

SSLv3

Enables SSL 3.0.

TLSv1d0

Enables TLS 1.0.

TLSv1d1

Enables TLS 1.1.

TLSv1d2

Enables TLS 1.2.

Guidelines

The **protocols** command specifies the SSL and TLS protocol versions to support.

request-client-auth

This command controls whether to request client authentication during the SSL handshake.

Syntax

request-client-auth { on | off }

Parameters

on

Requests client authentication.

off

Does not request client authentication. This setting is the default value.

Guidelines

The **request-client-auth** command controls whether to request client authentication during the SSL handshake.

When enabled, you must provide validation credentials with the **valcred** command to validate the SSL client certificate.

Related reference

[“valcred” on page 975](#)

This command specifies the name of the validation credentials to validate the SSL client certificate during the SSL handshake.

require-client-auth

This command controls whether to require client authentication during the SSL handshake.

Syntax

require-client-auth { on | off }

Parameters

on

Requires client authentication. This setting is the default value.

off

Does not require client authentication.

Guidelines

The **require-client-auth** command controls whether to require client authentication during the SSL handshake. When enabled, the SSL handshake fails if the client certificate is not provided.

This command is relevant only when the **request-client-auth** command is set to on.

send-client-auth-ca-list

This command controls whether to require client authentication during the SSL handshake.

Syntax

send-client-auth-ca-list { on | off }

Parameters

on

Transmits the client CA list during the SSL handshake. This setting is the default value.

off

Does not transmit the client CA list during the SSL handshake.

Guidelines

The **send-client-auth-ca-list** command controls whether to transmit the client CA list during the SSL handshake. Transmission of a client CA list is meaningful only when this profile uses validation credentials to validate the SSL client certificate.

Related reference

[“valcred” on page 975](#)

This command specifies the name of the validation credentials to validate the SSL client certificate during the SSL handshake.

ssl-options

This command specifies the list of elliptic curves that the SSL server profile supports.

Syntax

ssl-options *options*

Parameters

options

Specifies the options to apply to the SSL connection. The following values are valid. To specify multiple options, separate each option with the plus sign (+) character. For example, **max-duration** +**max-renegotiation**.

max-duration

Enables the option to specify the maximum duration of the SSL session.

max-renegotiation

Enables the option to specify the maximum number of the client initiated renegotiation that is allowed per session.

Guidelines

The **ssl-options** command specifies the options to apply to the SSL connection. Enabling these options has negative impact on the performance of the SSL communication. When enabled, you can change the following SSL settings. The default behavior is 60 seconds and 0 renegotiation attempts.

- Use the **max-duration** command to change the maximum duration of the SSL session.
- Use the **max-renegotiation-allowed** command to change the maximum number of renegotiation attempts that the client can initiate per session.

Related reference

[“max-duration” on page 971](#)

This command specifies the maximum time to maintain an SSL session after the initial negotiated handshake when the appliance is an SSL server.

valcred

This command specifies the name of the validation credentials to validate the SSL client certificate during the SSL handshake.

Syntax

valcred *valcred*

Parameters

valcred

Specifies the name of the validation credentials.

Guidelines

The **valcred** command specifies the name of the validation credentials to validate the SSL client certificate during the SSL handshake.

This command is required when the following conditions are true.

- The value set by the `request-client-auth` command is on.
- The value set by the `validate-client-cert` command is on.

Related reference

[“request-client-auth” on page 974](#)

This command controls whether to request client authentication during the SSL handshake.

[“validate-client-cert” on page 976](#)

This command controls whether to validate the client certificate during the SSL handshake if the client certificate is provided.

validate-client-cert

This command controls whether to validate the client certificate during the SSL handshake if the client certificate is provided.

Syntax

```
validate-client-cert { on | off }
```

Parameters

on

Validates the client certificate. This setting is the default value.

off

Does not validate the client certificate.

Guidelines

The **validate-client-cert** command controls whether to validate the client certificate during the SSL handshake if the client certificate is provided.

This command is relevant only when the **request-client-auth** command is set to on.

Related reference

[“request-client-auth” on page 974](#)

This command controls whether to request client authentication during the SSL handshake.

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.

<i>Table 60. System settings commands</i>	
Command name	Description
“admin-state” on page 708	This command sets the administrative state for the configuration.
“contact” on page 977	This command identifies the person or function responsible for appliance maintenance.
“custom-ui-file” on page 977	This command specifies the location of the custom user interface file.
“entitlement” on page 978	This command manages intrusion detection of the appliance in System Settings mode.
“locale” on page 978	This command specifies the locale for the operating language of the appliance.
“location” on page 979	This command specifies the location of the appliance.
“name” on page 979	This command sets an identifier for the appliance.

Related reference

[“system” on page 817](#)

This command enters System Settings mode.

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"
```

Related reference

[“location” on page 979](#)

This command specifies the location of the appliance.

[“name” on page 979](#)

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.

Related reference

[“language” on page 801](#)

This command enters the Language mode to set the administrative state of a language.

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 unique identifier.

- The name can contain a maximum of 127 ASCII alphanumeric characters and the following special characters: . - _
- The recommended length is up to 64 characters for compatibility with some remote systems.
- Some protocols do not support spaces. If the name contains spaces, enclose in double quotation marks.

Guidelines

The **name** command specifies the system identifier of the appliance.

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 269](#) and [“Changing appliance names in disaster recovery configurations” on page 290](#).

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 identifier.

```
# name DuLuth
```

- Set `Tango Lake` as the identifier.

```
# name "Tango Lake"
```

- Clear a previously set name:

```
# no name
```

Validation Credentials commands

Validation Credentials mode provides the commands to modify validation credentials.

To enter the mode, use the Crypto **valcred** command. To delete validation credentials, use the **no valcred** command.

While in this mode, use the following commands to define the validation credentials.

- 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.

Related reference

[“valcred” on page 749](#)

This command enters Validation Credentials mode.

cert-validation-mode

This command sets the certificate validation method.

Syntax

cert-validation-mode { *legacy* | *pkix* | *exact-match* }

Parameters

legacy

The validation credentials contain the exact peer certificate to match or the certificate of the immediate issuer, which might be an intermediate CA or a root CA. This mode is maintained for backwards compatibility. You can use *exact-match* or *pkix* in most cases instead of using *legacy*. This setting is the default value.

pkix

The complete certificate chain is checked from subject to root with this validation credentials for certificate validation. Validation succeeds only if the chain ends with a root certificate in the validation credentials. Nonroot certificates in the validation credentials are used as untrusted intermediate certificates. More untrusted intermediate certificates are dynamically obtained from the context at hand (SSL handshake messages, PKCS#7 tokens, PKIPath tokens, and so forth).

exact-match

The validation credentials contain the exact peer certificate to match. This mode is useful when you want to match the peer certificate exactly, but that certificate is not necessarily a self-signed (root) certificate.

Guidelines

The **cert-validation-mode** command sets the certificate validation method.

The *pkix* method, as described in RFC 3280, expects the remote peer to provide all intermediate certificates to the DataPower appliance during SSL negotiation. The associated validation credentials consist of self-signed certificates and certificates of trust anchors. Certificates can be a root CA or an intermediate CA.

Examples

- Create the ValCred-1 validation credentials with PKIX validation.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# cert-validation-mode pkix
```

- Restores the default setting for the ValCred-1 validation credentials.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# cert-validation-mode legacy
```

Related information

[RFC 3280, Internet X.509 Public Key Infrastructure: Certificate and Certificate Revocation List \(CRL\) Profile](#)

certificate

Adds a certificate alias to the validation credentials.

Syntax

certificate *alias*

no certificate *alias*

Parameters

alias

Specifies the name of the certificate alias.

Guidelines

The **certificate** command adds a certificate alias to the validation credentials.

A cryptographic profile optionally uses validation credentials to authenticate a remote SSL peer.

- Validation credentials are required by an SSL client only when it authenticates the certificate that is presented by the remote SSL server. Authentication of the server's certificate is not required by the SSL standard.
- Validation credentials are required by an SSL server only when it authenticates remote SSL clients. Authentication of SSL clients is not required by the SSL standard.

Assignment of validation credentials to a cryptographic profile requires 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.

Before you can add a certificate-alias to validation credentials, you must complete the following procedure.

1. Use the **copy** command to transfer the certificate to the appliance.
2. Use the Crypto **certificate** command to create an alias for the certificate.

Use the **no certificate** command to remove a certificate alias from the validation credentials.

Examples

Add the bob-1 certificate alias.

```
# certificate bob-1
```

Related reference

[“certificate” on page 727](#)

This command creates an alias for an X.509 certificate.

check-dates

This command controls whether to check the current date against the NotBefore value and the NotAfter value in the X.509 certificates and CRLs during certificate validation.

Syntax

check-dates { on | off }

Parameters

on

Checks the current date against the NotBefore value and the NotAfter value in the X.509 certificates and CRLs during certificate validation. If the certificate is expired, validation fails. This setting is the default value.

off

Does not check the current date against the NotBefore value and the NotAfter value in the X.509 certificates and CRLs during certificate validation. If a certificate is expired, it can still be accepted and validated during certificate validation without causing validation failure.

Guidelines

The **check-dates** command controls whether to check the current date and time against the activation date (the `NotBefore` value in the certificate) and the expiration date (the `NotAfter` value in the certificate) in the X.509 certificates and CRLs during certificate validation.

- When enabled, expired certificates causes the validation to fail during certificate validation.
- When disabled, expired certificates can be accepted and validated during certificate validation without causing the validation to fail.

Example

Allow the expired certificates to be accepted and validated during certificate validation.

```
# check-dates off
```

*crl*dp

This command controls support for the X.509 Certificate Distribution Point certificate extension.

Syntax

```
crldp{ ignore | require }
```

Parameters

ignore

Ignores the certificate extension. This setting is the default value.

require

Indicates that a candidate certificate is deemed valid if the presented certificate chain ends with a trust anchor. This method is used only when the validation credentials are for SSL peer validation.

Guidelines

The **crl`dp`** command controls support for the X.509 Certificate Distribution Point certificate extension. This noncritical certificate extension specifies how to obtain CRL information.

See RFC 2527 and RFC 3280 for information about certificate policies.

Examples

- Create the `ValCred-1` validation credentials that enable support the Certificate Distribution Point extension.

```
# valcred ValCred-1  
Crypto Validation Credentials configuration  
# crldp require  
#
```

- Restore the default state for the `ValCred-1` validation credentials.

```
Crypto Validation Credentials configuration  
# crldp ignore  
#
```

Related information

[RFC 2527, Internet X.509 Public Key Infrastructure: Certificate Policy and Certification Practices Framework](#)

[RFC 3280, Internet X.509 Public Key Infrastructure: Certificate and Certificate Revocation List \(CRL\) Profile](#)

explicit-policy

This command controls support for the initial explicit policy variable.

Syntax

explicit-policy

no explicit-policy

Guidelines

The **explicit-policy** command, controls support for the initial explicit policy variable. This command is meaningful only if the **cert-validation mode** command is set to `pkix`.

If enabled, the chain validation algorithm must end with a non-empty policy tree. If disabled, the algorithm can end with an empty policy tree unless policy constraints extensions in the chain require an explicit policy.

For information about certificate policies, see RFC 2527 and RFC 3280.

Examples

- Create the ValCred-1 validation credentials where the chain validation algorithm must end with an empty tree.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# cert-validation mode pkix
# explicit-policy
```

- Restore the default state for the ValCred-1 validation credentials.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# no explicit-policy
```

Related reference

[“cert-validation-mode” on page 981](#)

This command sets the certificate validation method.

Related information

[RFC 2527, Internet X.509 Public Key Infrastructure: Certificate Policy and Certification Practices Framework](#)

[RFC 3280, Internet X.509 Public Key Infrastructure: Certificate and Certificate Revocation List \(CRL\) Profile](#)

initial-policy-set

This command identifies a certificate policy for the validation credentials.

Syntax

initial-policy-set *identifier*

no initial-policy-set *identifier*

Parameters

identifier

Specifies the unique identifier for the certificate policy.

Guidelines

The **initial-policy-set** command, identifies a certificate policy for the validation credentials. This command is meaningful only if the **cert-validation mode** command is set to `pkix`.

RFC 2527 defines a certificate policy as follows:

A named set of rules that indicate the applicability of a certificate to a particular community and/or class of application with common security requirements. For example, a particular certificate policy might indicate applicability of a type of certificate to the authentication of electronic data interchange transactions for the trading of goods within a given price range.

The certificate policies extension contains a sequence of one or more policy information terms, each of which consists of an object identifier (OID) and optional qualifiers.

Note: The use of qualifiers is not supported. If present, they are ignored.

In a host certificate, policy information terms indicate the policy under which the certificate was issued and the purposes that the certificate can be used for.

In a CA certificate, policy information terms limit the set of policies for certification paths that include this certificate. When a CA does not want to limit the set of policies for certification paths that include this certificate, the CA can assert the special anyPolicy policy. The anyPolicy policy has an OID of 2.5.29.32.0.

Use this command as often as needed to construct a set of certificate policy identifiers. By default, the initial certificate policy set consists of the single OID 2.5.29.32.0, which identifies anyPolicy.

All members of the set are used during certificate chain processing as described in Section 6.1.1 of RFC 3280.

Use the **no initial-policy-set** command to remove a certificate policy from the validation credentials.

Examples

- Modify the ValCred-1 validation credentials to add the 1.3.6.1.4.1.14248.1.1 policy identifier.

```
# valcred ValCred-1
Crypto Validation Credentials configuration mode
# initial-policy-set 1.3.6.1.4.1.14248.1.1
```

- Modify the ValCred-1 validation credentials to remove the 1.3.6.1.4.1.14248.1.1 policy identifier.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# no initial-policy-set 1.3.6.1.4.1.14248.1.1
```

Related reference

[“cert-validation-mode” on page 981](#)

This command sets the certificate validation method.

[“explicit-policy” on page 984](#)

This command controls support for the initial explicit policy variable.

Related information

[RFC 2527, Internet X.509 Public Key Infrastructure: Certificate Policy and Certification Practices Framework](#)

[RFC 3280, Internet X.509 Public Key Infrastructure: Certificate and Certificate Revocation List \(CRL\) Profile](#)

require-crl

This command mandates CRL use during certificate chain processing.

Syntax

require-crl

no require-crl

Guidelines

The **require-crl** command mandates CRL (certificate revocation list) use during certificate chain processing. By default, CRL use is not required during certificate chains processing.

Use the **no require-crl** command to restore the default condition, which allows, but does not require, CRL use during certificate chains processing.

Examples

- Create the ValCred-1 validation credentials that require CRL use during certificate chain processing.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# require-crl
```

- Restores the default setting for the ValCred-1 validation credentials.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# no-require-crl
```

Related reference

[“use-crl” on page 986](#)

This command enables but does not require the use of certificate revocation lists during certificate chain processing.

use-crl

This command enables but does not require the use of certificate revocation lists during certificate chain processing.

Syntax

use-crl

no use-crl

Guidelines

The **use-crl** command enables but does not require the use of certificate revocation lists during certificate chain processing. By default, CRL use is enabled during processing certificate chains.

Use the **no use-crl** command to disable the use of certificate revocation lists during certificate chain processing.

Examples

- Create the ValCred-1 validation credentials that disable CRL use during certificate chain processing.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# no use-crl
```

- Restore the default setting for the ValCred-1 validation credentials.

```
# valcred ValCred-1
Crypto Validation Credentials configuration
# use-crl
```

Related reference

[“crl” on page 730](#)

This command enters CRL mode to create or modify a CRL update policy.

[“require-crl” on page 985](#)

This command mandates CRL use during certificate chain processing.

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
```

Related reference

[“over” on page 995](#)

This command sets the parent interface type.

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
```

Related reference

[“ip-config-mode” on page 989](#)

This command identifies the configuration mode for the VLAN interface.

[“ip-secondary-address” on page 991](#)

This command manages secondary network addresses for the VLAN interface.

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
```

Related reference

[“ip-address” on page 988](#)

This command assigns the primary network address for the VLAN interface.

[“ip-route” on page 990](#)

This command manages static routes in the routing table for the VLAN interface.

[“ip-secondary-address” on page 991](#)

This command manages secondary network addresses for the VLAN interface.

[“ipv4-default-gateway” on page 991](#)

This command designates the default IPv4 gateway for the VLAN interface.

[“ipv6-default-gateway” on page 992](#)

This command designates the default IPv6 gateway for the VLAN interface.

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
```

Related reference

[“ip-config-mode” on page 989](#)

This command identifies the configuration mode for the VLAN interface.

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.

Related reference

[“ip-config-mode” on page 989](#)

This command identifies the configuration mode for the VLAN interface.

[“ipv6-default-gateway” on page 992](#)

This command designates the default IPv6 gateway for the VLAN interface.

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.

Related reference

[“ipv6-nd-retransmit-timer” on page 993](#)

This command sets the interval between IPv6 neighbor discovery attempts for the VLAN interface.

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.

Related reference

[“ip-config-mode” on page 989](#)

This command identifies the configuration mode for the VLAN interface.

[“ipv4-default-gateway” on page 991](#)

This command designates the default IPv4 gateway for the VLAN interface.

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
```

Related reference

[“ethernet-interface” on page 987](#)

This command sets the Ethernet interface to provide connectivity.

[“link-aggregation-interface” on page 993](#)

This command sets the aggregate interface to provide connectivity.

[“over” on page 995](#)

This command sets the parent interface type.

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
```

Related reference

[“ethernet-interface” on page 987](#)

This command sets the Ethernet interface to provide connectivity.

[“link-aggregation-interface” on page 993](#)

This command sets the aggregate interface to provide connectivity.

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.

Guidelines

The `idle-timeout` command settings only apply to the IBM MQ Appliance web UI part of the user interface. The IBM MQ Console does not time out because it performs monitoring tasks.

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
```

Related reference

[“host-alias” on page 798](#)

This command enters Host Alias mode to map an IP address to an alias.

[“port” on page 998](#)

This command changes the listening port for the web management service.

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`.

REST management interface

The REST management interface provides access to the actions and to the configuration and status resources on the appliance.

The REST management interface is used to manage the appliance itself

For REST commands for administering IBM MQ, see [“Administering IBM MQ by using the REST API” on page 346](#), also see [Using the administrative REST API](#) in the IBM MQ documentation.

You can use the management REST interface to view status or configuration data, or to configure and reconfigure the appliance. You can send HTTP requests to the REST interface port and receive JSON-formatted responses with a payload and indication of success or failure. To send requests and receive responses, you can use the curl program, a similar shell tool, or a browser tool such as one for Firefox or Chrome. The response contains information about the resource at the URI target. You can also use the REST management interface by incorporating requests into programs that you write.

The REST interface uses a URI structure that makes the following resources available to work with:

- `/mgmt/config/default/configuration_objects`
- `/mgmt/status/default/status_objects`
- `/mgmt/actionqueue/default/`
- `/mgmt/filestore/`
- `/mgmt/metadata/`
- `/mgmt/types/`

REST management uses the HTTP GET, POST, PUT, and DELETE methods on its URIs. Not every resource is available for all HTTP methods. For example, status resources support only the GET method, where configuration resources allow more supported methods. You can retrieve the list of supported HTTP methods on any URI by sending the OPTIONS request to that URI.

The response payload conforms to the Hypertext Application Language standard. With this format, you can identify available resources in the JSON responses.

REST request structure

You can use different methods when you send requests to the REST management interface.

The general structure of all REST management interface requests is the same. You give the method (GET, PUT, POST, DELETE, OPTIONS) followed by the URI (starting with `https://address:5554/mgmt/...`). The remaining structure of the request depends on the resource and URI of the request. For example:

```
GET 'https://example.com:5554/mgmt/config/default/User/Bob'
```

See [“REST management resources summary” on page 1002](#) for valid URIs and applicable methods. You can also see which methods are supported for which URIs, by sending a request to that URI using the OPTIONS method. For example:

```
OPTIONS 'https://example.com:5554/mgmt/config/default/User/Bob'
```

When you make requests on config objects, you can optionally specify view, depth, and state query parameters. For example:

```
GET 'https://example.com:5554/mgmt/config/default/User/Bob?view=recursive&depth=2'
```

See [“REST management resources summary”](#) on page 1002 and [“Query parameters”](#) on page 1005 for details on the parameters and where you can use them.

Authentication header

An HTTP basic authentication header must be present in every request that is sent to the REST management interface.

Request payloads

When you use a PUT or POST method with a URI, you include a payload that contains what is to be put or posted.

The payload is in JSON and conforms to a specific schema. The payload schema is derived from the existing appliance SOMA schema, which is documented in `store:///xml-mgmt.xsd`.

For example, if you were wanted to modify the default gateway IP address in the eth0 Ethernet interface configuration, you would first retrieve the current value, for example:

```
GET https://myhost.com:5554/mgmt/config/default/EthernetInterface/eth0/DefaultGateway
```

```
{
  "_links" : {
    "self" : {
      "href" : "/mgmt/config/default/EthernetInterface/eth0/DefaultGateway"
    },
    "doc" : {
      "href" : "/mgmt/docs/config/EthernetInterface/DefaultGateway"
    }
  },
  "DefaultGateway" : "192.0.2.0"
}
```

You can then use a PUT request to modify the property:

```
PUT https://myhost.com:5554/mgmt/config/default/EthernetInterface/eth0/DefaultGateway
```

With the following payload:

```
{
  "DefaultGateway" : "192.0.2.56"
}
```

The appliance responds to the PUT request:

```
{
  "_links" : {
    "self" : { "href" : "/mgmt/config/default/EthernetInterface/eth0/DefaultGateway"
    },
    "doc" : { "href" : "/mgmt/docs/config/EthernetInterface/DefaultGateway"
    }
  },
  "DefaultGateway" : "property has been updated."
}
```

REST response structure

The appliance returns responses to REST management interface requests.

When you send a REST request, the appliance responds with a structured response in JSON format. The exact structure of the response depends on the resource and URI of the request, but all responses are similar.

The response includes all available resources from any point within the API. It includes the resources by embedding a `_links` element within the JSON response, which contains pointers to accessible

resources, including possible documentation. You can then target the resources within the `_links` element in your subsequent REST management interface requests.

For example, when you specify a GET request for the REST API root (`https://example.com:5554/mgmt/`), the response consists of the resources available that is structured as a list of fields and values:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/"
    }
  },
  "config": {
    "href": "/mgmt/config/"
  },
  "domains": {
    "href": "/mgmt/domains/config/"
  },
  "status": {
    "href": "/mgmt/status/"
  },
  "actionqueue": {
    "href": "/mgmt/actionqueue/"
  },
  "filestore": {
    "href": "/mgmt/filestore/"
  },
  "metadata": {
    "href": "/mgmt/metadata/"
  },
  "types": {
    "href": "/mgmt/types/"
  }
}
```

The following example shows the response to a request for status information about the current firmware version of the appliance, GET `'/mgmt/status/default/FirmwareVersion'`:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/FirmwareVersion"
    }
  },
  "doc": {
    "href": "/mgmt/docs/status/FirmwareVersion"
  }
},
"FirmwareVersion": {
  "Serial": "xxxxxxx",
  "Version": "MQ00.9.0.1.0",
  "Build": "xxxxxxx",
  "BuildDate": "2016/09/20 10:29:27",
  "WatchdogBuild": "MQ00.9.0.1.0",
  "InstalledDPOS": "MQ00.9.0.1.0",
  "RunningDPOS": "MQ00.9.0.1.0",
  "XMLAccelerator": "embedded",
  "MachineType": 5725,
  "ModelType": "S14"
}
}
```

REST management resources summary

You use URIs to work with appliance resources.

The root URI begins with the `/mgmt/` resource. All available resources are positioned below this resource.

The GET `https://address:5554/mgmt/` request returns a structure with the available resources.

Configuration resources

Use the configuration resources for the following actions:

- Viewing appliance configuration
- Configuring the appliance

Query parameters are available for use with these resources, see “Query parameters” on page 1005.

<i>Table 61. Configuration resources</i>			
URI	Supported HTTP methods	Supported query parameters	Description
/mgmt/config/	GET	none	List of appliance configurations
/mgmt/config/default/class	GET, POST	state	List of all configurations of a specific object class
/mgmt/config/default/class/object	GET, PUT, DELETE	depth, state, view	configuration of the specified instance in the object class
/mgmt/config/default/class/object/scalar_property	GET, PUT	state	The value of the specified scalar property belonging to the specified object in the class.
/mgmt/config/domain/class/object/vector_property	GET, PUT, POST	state	The value of the specified vector property belonging to the specified object in the class.

Status resources

Use the status resources to access appliance status information.

<i>Table 62. Status resources</i>		
URI	Supported HTTP methods	Description
/mgmt/status/	GET	List of all available status providers on the appliance
/mgmt/status/default/class	GET	The status provider of the specified object class

Actionqueue resources

Use the actionqueue resources to trigger operations on the appliance.

<i>Table 63. Actionqueue resources</i>		
URI	Supported HTTP methods	Description
/mgmt/actionqueue/	GET	Usage information for an actionqueue resource.
/mgmt/actionqueue/default/	POST	Run any action

Table 63. Actionqueue resources (continued)

URI	Supported HTTP methods	Description
/mgmt/actionqueue/default/operations	GET	List of all supported operations. The only actions that are supported are listed in the /mgmt/actionqueue/domain/operations URI. The actions that are excluded from the AnyActionElement type are specified in store:///xml-mgmt.xsd.
/mgmt/actionqueue/default/operations/operation	GET	JSON metadata information for the operation

Filestore resources

Use the filestore resources to work with files and directories on the appliance.

Table 64. Filestore resources

URI	Supported HTTP methods	Description
/mgmt/filestore/	GET	List of available URIs
/mgmt/filestore/default/	GET	List of available directories
/mgmt/filestore/default/directory	GET	Contents of the specified directory
/mgmt/filestore/default/directory/file	GET, PUT, POST, DELETE	Contents of the specified file

Metadata resources

Use these resources to access the metadata of appliance objects.

Table 65. Metadata resources

URI	Supported HTTP methods	Description
/mgmt/metadata/	GET	List of valid URIs for fetching metadata
/mgmt/metadata/default/class	GET	Complete metadata for a specific appliance configuration or status object in the specified object class
/mgmt/metadata/default/class/property	GET	Metadata for a specific property of an object class
/mgmt/metadata/default/operations/operation	GET, PUT, POST, DELETE	Metadata information for the specific actionqueue operation

Type resources

Use these resources to access object types.

Table 66. Type resources

URI	Supported HTTP methods	Description
/mgmt/types/	GET	List of all types in use on the appliance, including base types
/mgmt/types/default/type	GET	

Query parameters

You can use query parameters with requests to the configuration resource URIs.

You can access various kinds of data by using query parameters. All the parameters are optional. The parameters are used only on certain API requests, as specified. The default behavior is for the parameters not to be specified.

view

The view query parameter specifies the amount of detail that is returned in a request. Set `view=recursive` to receive the requested data and the contents of any referenced objects. Use with the depth parameter to control the depth of recursion (the default is seven levels of recursion). The recursive process can impact performance substantially.

- Each retrieved level of referenced configuration contents is a new depth. The depth parameter specifies how many levels to retrieve. The retrieved contents appear in the data under the `_embedded : descendants` field of the object.

```
"_embedded": {
  "descendants": [ ]
}
```

- The referenced data from all the depths is retrieved as a flat list of objects and types. No information about the hierarchical relationship between the returned data is conveyed by this list.

Example:

```
/mgmt/config/default/class/object?view=recursive&depth=2
```

depth

The depth query parameter limits the returned payload to a specified number of expansions, controlling the number of referenced objects that are retrieved. The value of the parameter is an integer in the range 1 - 7. The recursive process can impact performance substantially.

- This parameter is ignored unless it is used with the `view=recursive` parameter.
- When the `view=recursive` parameter is used without a depth parameter, the value of depth defaults to 7.

Example:

```
/mgmt/config/default/class/object?view=recursive&depth=4
```

state

The state query parameter retrieves the runtime state of a configuration object. The value of the state parameter can be any of `true | false, 1 | 0, yes | no`. The runtime state of an object is indicated by five labels:

- `opstate` indicates whether the object is operational. If an object is not operational, the `opstate` is down; otherwise, the `opstate` is up. For example, the `opstate` is down, when the object is incorrectly configured.

- `adminstate` indicates whether an object can be used. If an object can be used, the `adminstate` is enabled; otherwise, the `adminstate` is disabled.
- `eventcode` corresponds to the message ID associated with the current state of the object. When no error occurs, the `eventcode` is 0.
- `errorcode` contains the message text that is associated with the event code. When no error occurs, the `errorcode` is "".
- `configstate` indicates the saved state of the configuration. It can have one of the following values:
 - saved configuration is saved.
 - modified the configuration has been modified but not saved.
 - new the configuration is new but has not been saved.

Example:

```
/mgmt/config/default/class/object?state=true
```

The state of an object is retrieved with a GET request. The response contains both the runtime state and the configuration.

```
"state": {
  "opstate": "down",
  "adminstate": "disabled",
  "eventcode": "0x0034000d",
  "errorcode": "Object is disabled",
  "configstate": "saved"
}
```

REST API resources

This collection of topics provides reference information for each of the appliance REST API resources.

Status resources

The following URIs give access to the status resources on the appliance.

[/mgmt/status/](#)

You can use the HTTP GET method with the `/mgmt/status/` resource to list all available status providers on the appliance.

Note: You must include the trailing `'/'` in the URI.

GET

Use the HTTP GET method with the `/mgmt/status/` resource to request a list of appliance status providers.

Note: You must include the trailing `'/'` in the URI.

- [Resource URL](#)
- [“Request headers” on page 1007](#)
- [Request body format](#)
- [“Security requirements” on page 1007](#)
- [Response status codes](#)
- [“Response headers” on page 1007](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/
```

Where *port* is 5554 by default.

Request headers

None.

Request body format

None.

Security requirements

The caller does not need to be authenticated to make this call.

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes all status providers on the appliance. It includes the resources by embedding a `_links` element within the JSON response, which contains pointers to accessible status provider resources, including possible documentation. You can then target the resources within the `_links` element in your subsequent REST management interface requests. Each of these JSON objects contains the following attributes:

name

String.

Specifies the name of the status provider.

href

String.

Specifies the URI for the status provider

Examples

The following example requests a list of the status providers on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/
```

The following listing shows some fragments of the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/"
    },
    "ActiveUsers": {
```

```

    "href": "/mgmt/status/{domain}/ActiveUsers"
  },
  "Battery": {
    "href": "/mgmt/status/{domain}/CryptoModeStatus"
  },
  "ConnectionsAccepted": {
    "href": "/mgmt/status/{domain}/CurrentSensors"
  },
  ..
  "LogTargetStatus": {
    "href": "/mgmt/status/{domain}/LogTargetStatus"
  },
  "MQSystemResources": {
    "href": "/mgmt/status/{domain}/MQSystemResources"
  },
  "NDCacheStatus2": {
    "href": "/mgmt/status/{domain}/NDCacheStatus2"
  },
  ..
}
}

```

[/mgmt/status/default/ActiveUsers](#)

You can use the HTTP GET method with the `/mgmt/status/default/ActiveUsers` resource to list the currently active appliance users.

GET

Use the HTTP GET method with the `/mgmt/status/default/ActiveUsers` resource to request a list of current appliance users.

- [Resource URL](#)
- [“Request headers” on page 1008](#)
- [Request body format](#)
- [“Security requirements” on page 1008](#)
- [Response status codes](#)
- [“Response headers” on page 1009](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/ActiveUsers
```

Where `port` is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes all current users of the appliance within an "ActiveUsers" section. Each object contains the following fields:

session

Integer.

A system-assigned session identifier.

name

String.

The user name presented during login.

connection

String.

The method used to connect to the device, such as web-gui (using a browser), cli (through telnet or ssh), or serial (via a serial cable). The possible values are:

- serial
- ssh
- web-mgmt
- xml-mgmt
- system

address

String.

The IP address of the user's client machine. Not valid for serial connections.

login

String.

Time of login.

domain

String.

Always set to "default".

Examples

The following example requests a list of the current users on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/ActiveUsers
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/ActiveUsers"
    },
    "doc": {
      "href": "/mgmt/docs/status/ActiveUsers"
    }
  },
  "ActiveUsers": [
    {
      "session": 55168,
      "name": "admin",
      "connection": "web-gui",
      "address": "198.51.100.40",
      "login": "Mon Dec 24 10:21:02 2018",
      "domain": "default"
    },
    {
      "session": 55169,
      "name": "admin2",
      "connection": "saml-artifact",
      "address": "198.51.100.20",
      "login": "Mon Dec 24 10:23:09 2018",
      "domain": "default"
    }
  ]
}
```

[/mgmt/status/default/CryptoModeStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/CryptoModeStatus` resource to view the current crypto mode of the appliance (the crypto mode controls the encryption used for system management aspects of the appliance).

GET

Use the HTTP GET method with the `/mgmt/status/default/CryptoModeStatus` resource to request the current crypto mode setting.

- [Resource URL](#)
- [“Request headers” on page 1011](#)
- [Request body format](#)
- [“Security requirements” on page 1011](#)
- [Response status codes](#)
- [“Response headers” on page 1011](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/CryptoModeStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the current crypto mode settings. Each object contains the following fields:

TargetMode

String.

Indicates the cryptographic mode that the firmware is expected to be running in currently. This is the mode that is requested in the previous run of the firmware and this mode is expected to take effect at the previous firmware reload.

ActualMode

String.

Indicates the cryptographic mode that the firmware is actually running in. The current mode should be equivalent to the target mode unless a problem occurred during the previous firmware reload.

PendingTargetMode

String.

Indicates the cryptographic mode that the firmware is expected to run in after the next firmware reload.

Examples

The following example requests the current crypto mode setting of the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/CryptoModeStatus
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/CryptoModeStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/CryptoModeStatus"
    }
  },
  "CryptoModeStatus": {
    "TargetMode": "permissive",
    "ActualMode": "permissive",
    "PendingTargetMode": "permissive"
  }
}
```

Related reference

[“crypto-mode-set” on page 732](#)

This command sets the appliance-wide cryptographic mode for the next firmware reload.

[/mgmt/status/default/CurrentSensors](#)

You can use the HTTP GET method with the `/mgmt/status/default/CurrentSensors` resource to view status information about the power supply to the appliance. The status is derived from sensors that read electrical currents, from the power supplies and other components of the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/CurrentSensors` resource to request status information about the power supplied to the appliance, derived from sensors that read electrical currents, from the power supplies and other components of the appliance.

- [Resource URL](#)
- [“Request headers” on page 1012](#)
- [Request body format](#)
- [“Security requirements” on page 1013](#)
- [Response status codes](#)
- [“Response headers” on page 1013](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/CurrentSensors
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the status for the various voltages supplied to the appliance that is reported by the current sensors. Each object contains the following fields:

Name

String.

The name of the current sensor that is being monitored.

Value

Integer.

The most recent reading of the current sensor in milliamperes. There are only three significant digits.

UpperCriticalThreshold

Integer.

The maximum allowable reading of the current sensor in milliamperes.

ReadingStatus

String.

An indicator of whether the current reading is OK, or if the reading exceeds the upper critical threshold. The following status strings might be reported:

lowerNonRecoverable

The current sensor reading value is below the lower non-recoverable threshold, which indicates that there is a risk of damaging the hardware if the device is allowed to remain in this state. If you cannot diagnose and fix the cause of this reading, you should power down the device.

lowerCritical

The current sensor reading value is below the lower critical threshold, which indicates that the system is faulty, and might not be operating properly.

lowerNonCritical

The current sensor reading value is below the lower non-critical threshold, which is a warning that it is getting close to too low. The system should operate normally, but the cause of the sensor value should be investigated.

ok

The current sensor reading value is OK, and is within all limits.

upperNonCritical

The current sensor reading value is above the upper non-critical threshold, which is a warning that it is getting close to too high. The system should operate normally, but the cause of the sensor value should be investigated. This is commonly used on temperature sensors, where the investigation would be checking ambient air temperature and air flow.

upperCritical

The current sensor reading value is above the upper critical threshold, which indicates that the system is faulty, and might not be operating properly.

upperNonRecoverable

The current sensor reading value is above the upper non-recoverable threshold, which indicates that there is a risk of damaging the hardware if the device is allowed to remain in this state. If you cannot diagnose and fix the cause of this reading, you should power down the device.

failure

The reading indicates a failure. Used for non-numeric readings.

noReading

The attempt to read the sensor failed, so the status cannot be checked.

invalid

An internal software error has been detected checking the status.

Examples

The following example requests the power supply status of the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/CurrentSensors
```

The following listing shows the received response:

```
{
  "self": {
    "href": "/mgmt/status/default/CurrentSensors"
  },
  "doc": {
    "href": "/mgmt/docs/status/CurrentSensors"
  }
},
"CurrentSensors": [
  {
    "Name": "Current PU +12V 1",
    "Value": 8400,
    "UpperCriticalThreshold": 44800,
    "ReadingStatus": "ok"
  },
  {
    "Name": "Current PU +3.3V",
    "Value": 2000,
    "UpperCriticalThreshold": 16000,
    "ReadingStatus": "ok"
  },
  {
    "Name": "Current PU +5V",
    "Value": 2400,
    "UpperCriticalThreshold": 16000,
    "ReadingStatus": "ok"
  }
]
}
```

Related reference

[“crypto-mode-set” on page 732](#)

This command sets the appliance-wide cryptographic mode for the next firmware reload.

[/mgmt/status/default/DateTimeStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/DateTimeStatus` resource to request date and time settings and information from the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/DateTimeStatus` resource to request date and time settings and information from the appliance.

- [Resource URL](#)
- [“Request headers” on page 1015](#)
- [Request body format](#)
- [“Security requirements” on page 1015](#)
- [Response status codes](#)
- [“Response headers” on page 1016](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/DateTimeStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the data and time information for the appliance. Each object contains the following fields:

time

String.

The current local time.

timezone2

String.

The name of the time zone currently in effect.

tzspec

String.

The POSIX.1 timezone specification.

uptime2

String.

The total amount of time the appliance has been up since the last reload or reboot. Note that a shutdown and reload resets this counter.

bootuptime2

String.

The total amount of time since the last reboot. Note that this counter is reset by a shutdown and reboot, but not by a shutdown and reload.

Examples

The following example requests date and time information from the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/DateTimeStatus
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/DateTimeStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/DateTimeStatus"
    }
  },
  "DateTimeStatus": {
    "time": "Wed Jan 2 10:52:41 2019",
    "timezone": "GMT",
    "tzspec": "GMT0BST,M3.5.0/1:00,M10.5.0/2:00",
    "uptime2": "50 days 01:00:55",
    "bootuptime2": "50 days 01:01:41"
  }
}
```

Related concepts

[“Date, time, and locale configuration commands”](#) on page 750

You can use these commands to configure the time, date, and locale on the IBM MQ Appliance.

[/mgmt/status/default/DateTimeStatus2](#)

You can use the HTTP GET method with the `/mgmt/status/default/DateTimeStatus2` resource to request date and time settings and information from the appliance. (This call returns the same information as `/mgmt/status/default/DateTimeStatus`, but some of the returned information is formatted differently.)

GET

Use the HTTP GET method with the `/mgmt/status/default/DateTimeStatus2` resource to request date and time settings and information from the appliance.

- [Resource URL](#)
- [“Request headers” on page 1017](#)
- [Request body format](#)
- [“Security requirements” on page 1017](#)
- [Response status codes](#)
- [“Response headers” on page 1018](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/DateTimeStatus2
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the data and time information for the appliance. Each object contains the following fields:

time

String.

The current local time.

timezone

String.

The name of the time zone currently in effect.

tzspec

String.

The POSIX.1 timezone specification.

uptime3

String.

The total amount of time the appliance has been up since the last reload or reboot in ISO 8601 duration format. Note that a shutdown and reload resets this counter.

bootuptime3

String.

The total amount of time since the last reboot in ISO 8601 duration format. Note that this counter is reset by a shutdown and reboot, but not by a shutdown and reload.

Examples

The following example requests date and time information from the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/DateTimeStatus2
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/DateTimeStatus2"
    },
    "doc": {
      "href": "/mgmt/docs/status/DateTimeStatus2"
    }
  },
  "DateTimeStatus2": {
    "time": "Thu Jan 3 10:42:22 2019",
    "timezone": "GMT",
    "tzspec": "GMT0BST,M3.5.0/1:00,M10.5.0/2:00",
    "uptime3": "P51DT00H50M36S",
    "bootuptime3": "P51DT00H51M22S"
  }
}
```

Related concepts

[“Date, time, and locale configuration commands”](#) on page 750

You can use these commands to configure the time, date, and locale on the IBM MQ Appliance.

[/mgmt/status/default/DNSCacheHostStatus4](#)

You can use the HTTP GET method with the `/mgmt/status/default/DNSCacheHostStatus4` resource to request information about the host name/IP address pairs that are currently stored in the DNS cache. A new entry is stored each time the appliance successfully looks up a host name on a DNS server. The name can be either fully qualified (for example, `servername.uk.ibm.com`) or a partial host name (for example, `servername`).

GET

Use the HTTP GET method with the `/mgmt/status/default/DNSCacheHostStatus4` resource to request information about the host name/IP address pairs that are currently stored in the DNS cache.

- [Resource URL](#)
- [“Request headers” on page 1019](#)
- [Request body format](#)
- [“Security requirements” on page 1019](#)
- [Response status codes](#)
- [“Response headers” on page 1020](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/DNSCacheHostStatus4
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the server information in the DNS cache. Each object contains the following fields:

Hostname

String.

The host name that the appliance located on a configured DNS server using a configured DNS search domain.

IPVersion

String.

The version of the IP address: `ipv4` or `ipv6`.

IPAddress

String.

The numeric IP address corresponding to the given host name.

IPPref

String.

The preferred IP address version.

TTL

Integer.

The time to live that the DNS server provided for the host name.

Examples

The following example requests the server information in the DNS cache.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/DNSCacheHostStatus4
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/DNSCacheHostStatus4"
    },
    "doc": {
      "href": "/mgmt/docs/status/DNSCacheHostStatus4"
    }
  },
  "DNSCacheHostStatus4": [
    {
      "Hostname": "driftwood.hillesden.ibm.com",
      "IPVersion": "ipv4",
      "IPAddress": "192.0.2.170",
      "IPPref": "IPv4",
      "TTL": 21600
    },
    {
      "Hostname": "finlaggan.hillesden.ibm.com",
      "IPVersion": "ipv4",
      "IPAddress": "192.0.2.189",
    }
  ]
}
```

```

    "IPPref": "IPv4",
    "TTL": 21600
  },
  {
    "Hostname": "finlaggan.hillesden.ibm.com",
    "IPVersion": "ipv4",
    "IPAddress": "192.0.2.189",
    "IPPref": "IPv6",
    "TTL": 21600
  },
  {
    "Hostname": "timeserver.hillesden.ibm.com",
    "IPVersion": "ipv4",
    "IPAddress": "192.0.2.12",
    "IPPref": "IPv4",
    "TTL": 21600
  }
]
}

```

Related concepts

[“DNS commands” on page 764](#)

You can use the DNS commands to configure the DNS settings on the IBM MQ Appliance.

[/mgmt/status/default/DNSNameServerStatus2](#)

You can use the HTTP GET method with the `/mgmt/status/default/DNSNameServerStatus2` resource to request a list of configured and acquired name servers.

GET

Use the HTTP GET method with the `/mgmt/status/default/DNSNameServerStatus2` resource to request a list of configured and acquired name servers.

- [Resource URL](#)
- [“Request headers” on page 1021](#)
- [Request body format](#)
- [“Security requirements” on page 1022](#)
- [Response status codes](#)
- [“Response headers” on page 1022](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/DNSNameServerStatus2
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the name server information for the appliance. Each object contains the following fields:

SearchOrderIndex

Integer.

The name servers will be searched for in the specified order.

IPAddress

String.

The numeric IP address for any DNS server that is configured for use on this appliance.

UDPPort

Integer.

The UDP port number of the DNS server. (The default DNS UDP port is 53.) When the load balancing algorithm is first alive, this value is ignored.

TCPPort

Integer.

The TCP port number of the DNS server. (The default DNS TCP port is also 53.) When the load balancing algorithm is first alive, this value is ignored.

MaxRetries

Integer.

The maximum number of lookup requests that are sent to a given DNS name server before a connection failure is reported. However, if the load balancing algorithm is first alive, the individual host uses the top-level retry setting from DNS settings.

Examples

The following example requests the power supply status of the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/DNSNameServerStatus2
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/DNSNameServerStatus2"
    },
    "doc": {
      "href": "/mgmt/docs/status/DNSNameServerStatus2"
    }
  },
  "DNSNameServerStatus2": [
    {
      "SearchOrderIndex": 1,
      "IPAddress": "192.0.2.11",
      "UDPPort": 53,
      "TCPPort": 53,
      "MaxRetries": 3
    },
    {
      "SearchOrderIndex": 2,
      "IPAddress": "192.0.2.25",
      "UDPPort": 53,
      "TCPPort": 53,
      "MaxRetries": 3
    }
  ]
}
```

Related concepts

[“DNS commands” on page 764](#)

You can use the DNS commands to configure the DNS settings on the IBM MQ Appliance.

[/mgmt/status/default/DNSSearchDomainStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/DNSSearchDomainStatus` resource to request a list of configured and acquired search domains.

GET

Use the HTTP GET method with the `/mgmt/status/default/DNSSearchDomainStatus` resource to request a list of configured and acquired search domains.

- [Resource URL](#)
- [“Request headers” on page 1023](#)
- [Request body format](#)
- [“Security requirements” on page 1024](#)
- [Response status codes](#)
- [“Response headers” on page 1024](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/DNSSearchDomainStatus
```

Where `port` is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the search domain information for the appliance. Each object contains the following fields:

SearchDomain

String.

A list of configured and acquired domain names. The list is used if a partial host name is submitted to the DNS service. Each domain name is appended to the partial host name, and a DNS host name lookup is attempted for the `host.domain` pair until a name lookup succeeds or until the list is exhausted.

Examples

The following example requests a list of configured and acquired search domains on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/DNSSearchDomainStatus
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/DNSSearchDomainStatus"
    }
  },
}
```

```
"doc": {
  "href": "/mgmt/docs/status/DNSSearchDomainStatus"
},
"DNSSearchDomainStatus": [
  {
    "SearchDomain": "hursley.ibm.com"
  },
  {
    "SearchDomain": "v6.hursley.ibm.com"
  }
]
}
```

Related concepts

[“DNS commands” on page 764](#)

You can use the DNS commands to configure the DNS settings on the IBM MQ Appliance.

[/mgmt/status/default/DNSStaticHostStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/DNSStaticHostStatus` resource to request a list of configured static hosts.

GET

Use the HTTP GET method with the `/mgmt/status/default/DNSStaticHostStatus` resource to request a list of configured static hosts.

- [Resource URL](#)
- [“Request headers” on page 1025](#)
- [Request body format](#)
- [“Security requirements” on page 1025](#)
- [Response status codes](#)
- [“Response headers” on page 1026](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/DNSStaticHostStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the static host information for the appliance. Each object contains the following fields:

Hostname

String.

The host name that represents a DNS static host or a host alias.

IPAddress

String.

The IP address that corresponds to the given host name.

Type

String.

The type of configured host entry. Can be 'static' or 'Hostalias'.

Examples

The following example requests a list of configured static hosts on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/DNSStaticHostStatus
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/DNSStaticHostStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/DNSStaticHostStatus"
    }
  },
  "DNSStaticHostStatus": {
    "Hostname": "localhost",
    "IPAddress": "127.0.0.1",
    "Type": "static"
  }
}
```

Related concepts

[“DNS commands” on page 764](#)

You can use the DNS commands to configure the DNS settings on the IBM MQ Appliance.

[/mgmt/status/default/DomainStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/DomainStatus` resource to request the status of the default domain.

GET

Use the HTTP GET method with the `/mgmt/status/default/DomainStatus` resource to request the status of the default domain.

- [Resource URL](#)
- [“Request headers” on page 1027](#)
- [Request body format](#)
- [“Security requirements” on page 1027](#)
- [Response status codes](#)
- [“Response headers” on page 1028](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/DomainStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes the domain status for the appliance. Each object contains the following fields:

Domain

String.

The name of the domain. Only the default domain is available on the IBM MQ Appliance.

SaveNeeded

String.

Indicates whether the domain contains unsaved changes.

TraceEnabled

String.

Indicates whether the XML File Capture utility is enabled. This utility can only be enabled in the default domain.

DebugEnabled

String.

Indicates whether the domain is employing debug logging level. At this level many messages are written to the log, which impacts performance.

ProbeEnabled

String.

Indicates whether a service in the domain has the probe enabled.

DiagEnabled

String.

Indicates whether diagnostic tracing is enabled.

CurrentCommand

String.

Provides a list of commands that are currently running on the domain. Use this column to confirm whether long-running commands have completed. Possible values are:

- *no value* - there are no commands running on the domain.
- *no domain* - an Application Domain object is being deleted.
- *restart domain* - the domain is being restored to a state defined in a saved configuration file.
- *reset domain* - objects in the domain configuration are being deleted.
- *write memswitch* - the configuration is being saved.
- *import-execute* - importing an Import Package object.
- *exec* - a configuration script is running.
- *domain* - the domain is being created.
- *backup* - an archive of configuration data is being generated.
- *save chkpoint* - a checkpoint configuration file is being generated.
- *test schema* - an XML file is being tested against a schema.
- *umount* - a file system is being unmounted.
- *configuring* - the domain is being configured.

- secure-backup - the domain is being securely backed up.
- secure-restore - the domain is being restored from a secure backup.

QuiesceState

String.

This setting does not apply to the default domain.

InterfaceState

String.

Indicates the state of the networking interfaces of the appliance. The value gives an indication as to which networking objects require attention. Possible values are:

- Appliance OK
Indicates that all networking is up on the appliance.
- Link errors
Indicates that one or more ethernet links are misconfigured on the appliance.
- VLAN errors
Indicates that one or more VLANs are misconfigured on the appliance.
- Aggregation errors
Indicates that one or more Aggregations are misconfigured on the appliance.
- Link/VLAN errors
Indicates that one or more ethernet links and VLANs are misconfigured on the appliance.
- Link/Aggregation errors
Indicates that one or more ethernet links and Aggregations are misconfigured on the appliance.
- VLAN/Aggregation errors
Indicates that one or more VLANs and Aggregations are misconfigured on the appliance.
- All errors
Indicates that one or more ethernet links, VLANs, and Aggregations are misconfigured on the appliance.

FailsafeMode

String.

Indicates the failsafe startup mode of the appliance. Failsafe modes are entered after the appliance unexpectedly restarts one or more times. Until you correct the problem, the appliance cannot accept and process client requests. Contact IBM support. Possible values are:

- **none**
The appliance is not in a failsafe mode and is operating normally.
- **management**
Only the default domain and configured Telnet, SSH and web management interfaces are available.
- **console**
Only the default domain and a serial console interface are available.

Examples

The following example requests the status of the default domain on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/DomainStatus
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/DomainStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/DomainStatus"
    }
  },
  "DomainStatus": {
    "Domain": "default",
    "SaveNeeded": "off",
    "TraceEnabled": "off",
    "DebugEnabled": "off",
    "ProbeEnabled": "off",
    "DiagEnabled": "off",
    "CurrentCommand": "",
    "QuiesceState": "",
    "InterfaceState": "ok",
    "FailsafeMode": "none"
  }
}
```

[/mgmt/status/default/EnvironmentalFanSensors](#)

You can use the HTTP GET method with the `/mgmt/status/default/EnvironmentalFanSensors` resource to request information from fan sensors on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/EnvironmentalFanSensors` resource to request the information from the sensors monitoring the fans on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1030](#)
- [Request body format](#)
- [“Security requirements” on page 1031](#)
- [Response status codes](#)
- [“Response headers” on page 1031](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/EnvironmentalFanSensors
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information from the fan sensors and contains the following fields:

FanID

String.

The name of the fan.

FanSpeed

Integer.

The speed of the fan in revolutions per minute (RPM).

LowerCriticalThreshold

Integer.

The lowest allowable reading of the fan speed sensor.

ReadingStatus

String.

An indicator of whether the current fan speed is OK or the fan speed is below the lower critical threshold. If the status is not OK, replace the fan.

- **lowerNonRecoverable**

The current sensor reading value is below the lower non-recoverable threshold, which indicates that there is a risk of damaging the hardware if the device is allowed to remain in this state. If you cannot alleviate the cause of this reading, you should power down the device.

- **lowerCritical**

The current sensor reading value is below the lower critical threshold, which indicates that the system is faulty, and may not be operating properly.

- **lowerNonCritical**

The current sensor reading value is below the lower non-critical threshold, which is a warning that it is getting close to too low. The system should operate normally, but the cause of the sensor value should be investigated.

- **ok**

The current sensor reading value is OK, and is within all limits.

- **upperNonCritical**

The current sensor reading value is above the upper non-critical threshold, which is a warning that it is getting close to too high. The system should operate normally, but the cause of the sensor value should be investigated. This is commonly used on temperature sensors, where the investigation would be checking ambient air temperature and air flow.

- **upperCritical**

The current sensor reading value is above the upper critical threshold, which indicates that the system is faulty, and may not be operating properly.

- **upperNonRecoverable**

The current sensor reading value is above the upper non-recoverable threshold, which indicates that there is a risk of damaging the hardware if the device is allowed to remain in this state. If you cannot alleviate the cause of this reading, you should power down the device.

- **failure**

The reading indicates a failure. Used for non-numeric readings.

- **noReading**

The attempt to read the sensor failed, so the status cannot be checked.

- **invalid**

An internal software error has been detected checking the status.

Examples

The following example requests the status of the appliance fans.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/EnvironmentalFanSensors
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/EnvironmentalFanSensors"
    },
    "doc": {
      "href": "/mgmt/docs/status/EnvironmentalFanSensors"
    }
  },
  "EnvironmentalFanSensors": [
    {
      "FanID": "fan-1",
      "FanSpeed": 3900,
      "LowerCriticalThreshold": 1200,
      "ReadingStatus": "ok"
    },
    {
      "FanID": "fan-2",
      "FanSpeed": 3900,
      "LowerCriticalThreshold": 1200,
      "ReadingStatus": "ok"
    },
    {
      "FanID": "fan-3",
      "FanSpeed": 3900,
      "LowerCriticalThreshold": 1200,
      "ReadingStatus": "ok"
    }
  ]
}
```

```
}  
  ]  
}
```

[/mgmt/status/default/EnvironmentalSensors](#)

You can use the HTTP GET method with the `/mgmt/status/default/EnvironmentalSensors` resource to request information from various sensors on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/EnvironmentalSensors` resource to request the information from sensors on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1033](#)
- [Request body format](#)
- [“Security requirements” on page 1033](#)
- [Response status codes](#)
- [“Response headers” on page 1034](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/EnvironmentalSensors
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information from the fan sensors. Each object contains the following fields:

systemTemp

Integer.

Ambient temperature.

cpu1Temp

Integer.

CPU 1 temperature

cpu2Temp

Integer.

CPU 2 temperature

cpu1rpm

Integer.

CPU 1 fan speed

cpu2rpm

Integer.

CPU 2 fan speed

chassis1rpm

Integer.

Chassis fan 1 speed

chassis2rpm

Integer.

Chassis fan 2 speed

chassis3rpm

Integer.

Chassis fan 3 speed

caseopen

String.

Intrusion detected

volt33

String.

3.3 voltage

volt5

String.

5 voltage

volt12

String.

12 voltage

powersupply

String.

Power supply status

batteryinstalldate

String.

Battery installation date

checkintrusion

String.

Check intrusion detection

Examples

The following example requests the status of the appliance fans.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/EnvironmentalSensors
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/EnvironmentalSensors"
    },
    "doc": {
      "href": "/mgmt/docs/status/EnvironmentalSensors"
    }
  },
  "EnvironmentalSensors": {
    "systemTemp": 29,
    "cpu1Temp": 31,
    "cpu2Temp": 37,
    "cpu1rpm": 0,
    "cpu2rpm": 0,
    "chassis1rpm": 0,
    "chassis2rpm": 0,
    "chassis3rpm": 0,
    "caseopen": "no",
    "volt33": "3.31",
    "volt5": "5.06",
    "volt12": "11.9",
    "powersupply": "power-1-failure",
    "batteryinstalldate": "Unavailable",
    "checkintrusion": "yes"
  }
}
```

[/mgmt/status/default/EthernetCountersStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/EthernetCountersStatus` resource to request detailed statistics for each Ethernet device on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/EthernetCountersStatus` resource to request detailed statistics for each Ethernet device on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1036](#)
- [Request body format](#)
- [“Security requirements” on page 1036](#)
- [Response status codes](#)
- [“Response headers” on page 1036](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/EthernetCountersStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information from the Ethernet interfaces. Each object contains the following fields:

InterfaceIndex

Integer.

The value of the SNMP `ifIndex` column in the `ifTable` for this interface.

Name

String.

The name of the ethernet interface.

InUnicastPackets

Integer.

The number of unicast packets received on this interface.

InMulticastPackets

Integer.

The number of multicast packets received on this interface.

InBroadcastPackets

Integer.

The number of broadcast packets received on this interface.

OutUnicastPackets

String.

The number of unicast packets transmitted on this interface.

OutMulticastPackets

Integer.

The number of multicast packets transmitted on this interface.

OutBroadcastPackets

Integer.

The number of broadcast packets transmitted on this interface.

InOctets

Integer.

The number of bytes received on this interface at the MAC level.

OutOctets

Integer.

The number of bytes transmitted on this interface at the MAC level.

InErrors

Integer.

The total number of receive errors on this interface.

OutErrors

Integer.

The total number of transmit errors on this interface.

OutDiscards

Integer.

The number of packets not transmitted for flow control reasons.

AlignmentErrors

Integer.

The number of packets received on this interface that were not an integral number of bytes in length.

FCSErrors

Integer.

The number of packets received on this interface with an invalid Frame Check Sequence (checksum). This does not include FCS errors on packets that were too long or too short.

SingleCollisionFrames

Integer.

The number of packets successfully transmitted on this interface after a single collision. This can only happen when the interface is running in Half-Duplex mode.

MultipleCollisionFrames

Integer.

The number of packets successfully transmitted on this interface after multiple collisions, but less than 16 collisions. This can only happen when the interface is running in Half-Duplex mode.

SQETestErrors

Integer.

The number of times that an SQE test error was encountered. This only can happen when the link is operating in 10BASE-T Half-Duplex.

DeferredTransmissions

Integer.

The number of frames for which the first transmission attempt was deferred because the medium was busy. This can only happen when the interface is running in Half-Duplex mode.

LateCollisions

Integer.

The number of times that a collision was detected later than one slot time after the transmission of a packet. This can only happen when the interface is running in Half-Duplex mode.

ExcessiveCollisions

Integer.

The number of times that transmission of a packet failed because it encountered sixteen collisions in a row. This can only happen when the interface is running in Half-Duplex mode.

InternalMacTransmitErrors

String.

The number of times that transmission of packets failed due to errors inside the MAC layer of the Ethernet interface. These may be due to temporary resource limitations.

CarrierSenseErrors

Integer.

The number transmitted packets during which there were failures of carrier sense. This can only happen when the interface is running in Half-Duplex mode.

FrameTooShorts

Integer.

The number of received packets that were shorter than 64 bytes. This can be the result of a collision. Such packages are also known as runts.

FrameTooLongs

Integer.

The number of received packets that were longer than the configured MTU. This can be the result of a collision, as well as due to incompatible configuration.

InternalMacReceiveErrors

Integer.

The number of times that reception of packets failed due to errors inside the MAC layer of the Ethernet interface. These may be due to temporary resource limitations.

InPauseFrames

Integer.

The number of pause frames received.

OutPauseFrames

Integer.

The number of pause frames transmitted.

Examples

The following example requests detailed statistics for each Ethernet device.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/EthernetCountersStatus
```

The following listing shows sections from the received response:

```
{  
  "_links": {
```

```

"self": {
  "href": "/mgmt/status/default/EthernetCountersStatus"
},
"doc": {
  "href": "/mgmt/docs/status/EthernetCountersStatus"
}
},
"EthernetCountersStatus": [
  {
    "InterfaceIndex": 4,
    "Name": "mgt0",
    "InUnicastPackets": "",
    "InMulticastPackets": 25935923,
    "InBroadcastPackets": 1805817,
    "OutUnicastPackets": "",
    "OutMulticastPackets": 6,
    "OutBroadcastPackets": 59,
    "InOctets": 2361799647,
    "OutOctets": 682384283,
    "InErrors": 0,
    "OutErrors": 0,
    "OutDiscards": 0,
    "AlignmentErrors": 0,
    "FCSerrors": 0,
    "SingleCollisionFrames": 0,
    "MultipleCollisionFrames": 0,
    "SQETestErrors": 0,
    "DeferredTransmissions": 0,
    "LateCollisions": 0,
    "ExcessiveCollisions": 0,
    "InternalMacTransmitErrors": "",
    "CarrierSenseErrors": 0,
    "FrameTooShorts": 0,
    "FrameTooLongs": 0,
    "InternalMacReceiveErrors": 0,
    "InPauseFrames": 0,
    "OutPauseFrames": 0
  },
  .....
  {
    "InterfaceIndex": 15,
    "Name": "eth17",
    "InUnicastPackets": "",
    "InMulticastPackets": 0,
    "InBroadcastPackets": 0,
    "OutUnicastPackets": "",
    "OutMulticastPackets": 0,
    "OutBroadcastPackets": 0,
    "InOctets": 0,
    "OutOctets": 0,
    "InErrors": 0,
    "OutErrors": 0,
    "OutDiscards": 0,
    "AlignmentErrors": 0,
    "FCSerrors": 0,
    "SingleCollisionFrames": 0,
    "MultipleCollisionFrames": 0,
    "SQETestErrors": 0,
    "DeferredTransmissions": 0,
    "LateCollisions": 0,
    "ExcessiveCollisions": 0,
    "InternalMacTransmitErrors": "",
    "CarrierSenseErrors": 0,
    "FrameTooShorts": 0,
    "FrameTooLongs": 0,
    "InternalMacReceiveErrors": 0,
    "InPauseFrames": 0,
    "OutPauseFrames": 0
  }
]
}

```

[/mgmt/status/default/EthernetMAUStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/EthernetMAUStatus` resource to request the device-specific status of Ethernet MAUs on the system..

GET

Use the HTTP GET method with the `/mgmt/status/default/EthernetMAUStatus` resource to request the device-specific status of Ethernet MAUs on the system.

- [Resource URL](#)
- [“Request headers” on page 1040](#)
- [Request body format](#)
- [“Security requirements” on page 1040](#)
- [Response status codes](#)
- [“Response headers” on page 1041](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/EthernetMAUStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information from the Ethernet MAUs. Each object contains the following fields:

InterfaceIndex

Integer.

The value of the SNMP `ifIndex` column in the `ifTable` for this interface.

Name

String.

The name of the ethernet interface.

MauIndex

Integer.

The index of this particular MAU on this Ethernet interface.

Status

String.

The current MAU link status:

- **other** other
- **unknown** unknown at present time
- **available** Available, normal
- **notAvailable** Not available
- **remoteFault** Down due to remote fault
- **offline** Offline

JabberState

Integer.

The current MAU jabber error state:

- **other** other
- **unknown** unknown at present time
- **noJabber** No jabber, normal
- **jabbering** Jabber error

AutoNegRemoteSignaling

String.

Does the remote end support auto-negotiation. Will be **notDetected** if no FLP bursts received.

- **detected** Detected, FLP bursts seen
- **notDetected** Not Detected, no FLP bursts seen

AutoNegConfig

String.

The current status of the auto-negotiation process.

- **other** Other
- **configuring** Configuring, auto-negotiation in progress
- **complete** Complete
- **disabled** Disabled
- **parallelDetectFail** Parallel detection failed

AutoNegCapabilities

String.

The capabilities of this Ethernet MAU:

- **eth10baseTHD** 10BASE-T half-duplex
- **eth10baseTFD** 10BASE-T full-duplex
- **eth100baseXHD** 10BASE-X half-duplex
- **eth100baseXFD** 10BASE-X full-duplex
- **eth100baseT4HD** 10BASE-T4 half-duplex
- **eth100baseTXHD** 10BASE-TX half-duplex
- **eth100baseTXFD** 10BASE-TX full-duplex
- **eth100baseT2HD** 10BASE-T2 half-duplex
- **eth100baseT2FD** 10BASE-T2 full-duplex
- **eth1000baseTHD** 1000BASE-T half-duplex
- **eth1000baseTFD** 1000BASE-T full-duplex
- **eth1000baseXHD** 1000BASE-X half-duplex
- **eth1000baseXFD** 1000BASE-X full-duplex
- **ethAsymPauseFD** Asymmetric PAUSE toward link partner for full-duplex
- **ethSymPauseFD** Symmetric PAUSE for full-duplex
- **ethAsymSymPauseFD** Asymmetric and symmetric PAUSE toward local device for full-duplex

AutoNegCapabilitiesAdvertised

String.

The capabilities of this Ethernet MAU being advertised in auto-negotiation.

- **eth10baseTHD** 10BASE-T half-duplex
- **eth10baseTFD** 10BASE-T full-duplex
- **eth100baseXHD** 10BASE-X half-duplex
- **eth100baseXFD** 10BASE-X full-duplex
- **eth100baseT4HD** 10BASE-T4 half-duplex
- **eth100baseTXHD** 10BASE-TX half-duplex
- **eth100baseTXFD** 10BASE-TX full-duplex
- **eth100baseT2HD** 10BASE-T2 half-duplex
- **eth100baseT2FD** 10BASE-T2 full-duplex
- **eth1000baseTHD** 1000BASE-T half-duplex
- **eth1000baseTFD** 1000BASE-T full-duplex
- **eth1000baseXHD** 1000BASE-X half-duplex
- **eth1000baseXFD** 1000BASE-X full-duplex
- **ethAsymPauseFD** Asymmetric PAUSE toward link partner for full-duplex
- **ethSymPauseFD** Symmetric PAUSE for full-duplex
- **ethAsymSymPauseFD** Asymmetric and symmetric PAUSE toward local device for full-duplex

AutoNegCapabilitiesLinkPartner

String.

The capabilities of the peer Ethernet MAU received by auto-negotiation.

- **eth10baseTHD** 10BASE-T half-duplex
- **eth10baseTFD** 10BASE-T full-duplex

- **eth100baseXHD** 10BASE-X half-duplex
- **eth100baseXFD** 10BASE-X full-duplex
- **eth100baseT4HD** 10BASE-T4 half-duplex
- **eth100baseTXHD** 10BASE-TX half-duplex
- **eth100baseTXFD** 10BASE-TX full-duplex
- **eth100baseT2HD** 10BASE-T2 half-duplex
- **eth100baseT2FD** 10BASE-T2 full-duplex
- **eth1000baseTHD** 1000BASE-T half-duplex
- **eth1000baseTFD** 1000BASE-T full-duplex
- **eth1000baseXHD** 1000BASE-X half-duplex
- **eth1000baseXFD** 1000BASE-X full-duplex
- **ethAsymPauseFD** Asymmetric PAUSE toward link partner for full-duplex
- **ethSymPauseFD** Symmetric PAUSE for full-duplex
- **ethAsymSymPauseFD** Asymmetric and symmetric PAUSE toward local device for full-duplex

AutoNegRemoteFaultAdvertised

String.

Whether a local fault is causing advertising of remote fault to peer:

- **noError** No error
- **offline** Offline
- **linkFailure** Link failure
- **autoNegError** Auto-negotiation failure

AutoNegRemoteFaultReceived

String.

Whether a remote fault is causing reception of remote fault from peer:

- **noError** No error
- **offline** Offline
- **linkFailure** Link failure
- **autoNegError** Auto-negotiation failure

Examples

The following example requests the device-specific status of Ethernet MAUs on the system.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/EthernetMAUStatus
```

The following listing shows sections from the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/EthernetMAUStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/EthernetMAUStatus"
    }
  },
  "EthernetMAUStatus": [
    {
      "InterfaceIndex": 4,
      "Name": "mgt0",
      "MauIndex": 1,
      "Status": "available",
```

```

"JabberState": "noJabber",
"AutoNegRemoteSignaling": "detected",
"AutoNegConfig": "complete",
"AutoNegCapabilities": {
  "eth10baseTHD": "on",
  "eth10baseTFD": "on",
  "eth100baseXHD": "on",
  "eth100baseXFD": "on",
  "eth100baseT4HD": "off",
  "eth100baseTXHD": "off",
  "eth100baseTXFD": "off",
  "eth100baseT2HD": "off",
  "eth100baseT2FD": "off",
  "eth1000baseTHD": "on",
  "eth1000baseTFD": "on",
  "eth1000baseXHD": "off",
  "eth1000baseXFD": "off",
  "ethAsymPauseFD": "off",
  "ethSymPauseFD": "off",
  "ethAsymSymPauseFD": "off"
},
"AutoNegCapabilitiesAdvertised": {
  "eth10baseTHD": "on",
  "eth10baseTFD": "on",
  "eth100baseXHD": "off",
  "eth100baseXFD": "off",
  "eth100baseT4HD": "off",
  "eth100baseTXHD": "on",
  "eth100baseTXFD": "on",
  "eth100baseT2HD": "off",
  "eth100baseT2FD": "off",
  "eth1000baseTHD": "on",
  "eth1000baseTFD": "on",
  "eth1000baseXHD": "off",
  "eth1000baseXFD": "off",
  "ethAsymPauseFD": "off",
  "ethSymPauseFD": "off",
  "ethAsymSymPauseFD": "on"
},
"AutoNegCapabilitiesLinkPartner": {
  "eth10baseTHD": "on",
  "eth10baseTFD": "on",
  "eth100baseXHD": "off",
  "eth100baseXFD": "off",
  "eth100baseT4HD": "off",
  "eth100baseTXHD": "on",
  "eth100baseTXFD": "on",
  "eth100baseT2HD": "off",
  "eth100baseT2FD": "off",
  "eth1000baseTHD": "off",
  "eth1000baseTFD": "on",
  "eth1000baseXHD": "off",
  "eth1000baseXFD": "off",
  "ethAsymPauseFD": "off",
  "ethSymPauseFD": "on",
  "ethAsymSymPauseFD": "off"
},
"AutoNegRemoteFaultAdvertised": "noError",
"AutoNegRemoteFaultReceived": "noError"
},
.....
{
  "InterfaceIndex": 15,
  "Name": "eth17",
  "MauIndex": 1,
  "Status": "notAvailable",
  "JabberState": "noJabber",
  "AutoNegRemoteSignaling": "notDetected",
  "AutoNegConfig": "configuring",
  "AutoNegCapabilities": {
    "eth10baseTHD": "on",
    "eth10baseTFD": "on",
    "eth100baseXHD": "on",
    "eth100baseXFD": "on",
    "eth100baseT4HD": "off",
    "eth100baseTXHD": "off",
    "eth100baseTXFD": "off",
    "eth100baseT2HD": "off",
    "eth100baseT2FD": "off",
    "eth1000baseTHD": "on",
    "eth1000baseTFD": "on",
    "eth1000baseXHD": "off",

```

```

    "eth1000baseXFD": "off",
    "ethAsymPauseFD": "off",
    "ethSymPauseFD": "off",
    "ethAsymSymPauseFD": "off"
  },
  "AutoNegCapabilitiesAdvertised": {
    "eth10baseTHD": "on",
    "eth10baseTFD": "on",
    "eth100baseXHD": "off",
    "eth100baseXFD": "off",
    "eth100baseT4HD": "off",
    "eth100baseTXHD": "on",
    "eth100baseTXFD": "on",
    "eth100baseT2HD": "off",
    "eth100baseT2FD": "off",
    "eth1000baseTHD": "off",
    "eth1000baseTFD": "on",
    "eth1000baseXHD": "off",
    "eth1000baseXFD": "off",
    "ethAsymPauseFD": "off",
    "ethSymPauseFD": "off",
    "ethAsymSymPauseFD": "on"
  },
  "AutoNegCapabilitiesLinkPartner": {
    "eth10baseTHD": "off",
    "eth10baseTFD": "off",
    "eth100baseXHD": "off",
    "eth100baseXFD": "off",
    "eth100baseT4HD": "off",
    "eth100baseTXHD": "off",
    "eth100baseTXFD": "off",
    "eth100baseT2HD": "off",
    "eth100baseT2FD": "off",
    "eth1000baseTHD": "off",
    "eth1000baseTFD": "off",
    "eth1000baseXHD": "off",
    "eth1000baseXFD": "off",
    "ethAsymPauseFD": "off",
    "ethSymPauseFD": "off",
    "ethAsymSymPauseFD": "off"
  },
  "AutoNegRemoteFaultAdvertised": "noError",
  "AutoNegRemoteFaultReceived": "noError"
}
]
}

```

[/mgmt/status/default/EthernetMIISatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/EthernetMIISatus` resource to request detailed statistics available for each Ethernet MII registers on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/EthernetMIISatus` resource to request detailed statistics available for Ethernet MII registers on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1046](#)
- [Request body format](#)
- [“Security requirements” on page 1046](#)
- [Response status codes](#)
- [“Response headers” on page 1046](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/EthernetMIISatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information from the Ethernet MII registers. Each object contains the following fields:

InterfaceIndex

Integer.

The value of the SNMP `ifIndex` column in the `ifTable` for this interface.

Name

String.

The name of the ethernet interface.

MauIndex

Integer.

The index of this particular MAU on this Ethernet interface.

Control

String.

The current read value of the MII control register on this MAU.

Status

String.

The current read value of the MII status register on this MAU.

PhyIdentifierHigh

String.

The current read value of the MII PHY identifier high register on this MAU.

PhyIdentifierLow

String.

The current read value of the MII PHY identifier low register on this MAU.

AutoNegAdvertisement

String.

The current read value of the MII auto-negotiation advertisement register on this MAU.

AutoNegLinkPartnerAbility

String.

The current read value of the MII auto-negotiation link partner base page ability register on this MAU.

AutoNegExpansion

String.

The current read value of the MII auto-negotiation expansion register on this MAU.

AutoNegNextPageTransmit

String.

The current read value of the MII auto-negotiation next page transmit register on this MAU.

AutoNegLinkPartnerRecNextPage

String.

The current read value of the MII auto-negotiation link partner received next page register on this MAU.

MasterSlaveControlRegister

String.

The current read value of the MII master slave control register on this MAU.

MasterSlaveStatusRegister

String.

The current read value of the MII master slave status register on this MAU.

PSEControl

String.

The current read value of the MII PSE control register on this MAU.

PSEStatus

String.

The current read value of the MII PSE status register on this MAU.

MMDAccessControl

String.

The current read value of the MII MMD access control register on this MAU.

MMDAccessAddressData

String.

The current read value of the MII MMD access address data register on this MAU.

ExtendedStatus

String.

The current read value of the MII extended status register on this MAU.

Vendor16

String.

The current read value of the MII vendor register 16 on this MAU.

Vendor17

String.

The current read value of the MII vendor register 17 on this MAU.

Vendor18

String.

The current read value of the MII vendor register 18 on this MAU.

Vendor19

String.

The current read value of the MII vendor register 19 on this MAU.

Vendor20

String.

The current read value of the MII vendor register 20 on this MAU.

Vendor21

String.

The current read value of the MII vendor register 21 on this MAU.

Vendor22

String.

The current read value of the MII vendor register 22 on this MAU.

Vendor23

String.

The current read value of the MII vendor register 23 on this MAU.

Vendor24

String.

The current read value of the MII vendor register 24 on this MAU.

Vendor25

String.

The current read value of the MII vendor register 25 on this MAU.

Vendor26

String.

The current read value of the MII vendor register 26 on this MAU.

Vendor27

String.

The current read value of the MII vendor register 27 on this MAU.

Vendor28

String.

The current read value of the MII vendor register 28 on this MAU.

Vendor29

String.

The current read value of the MII vendor register 29 on this MAU.

Vendor30

String.

The current read value of the MII vendor register 30 on this MAU.

Vendor31

String.

The current read value of the MII vendor register 31 on this MAU.

Examples

The following example requests the status of the ethernet MII registers.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/EthernetMIIStatus
```

The following listing shows sections from the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/EthernetMIIStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/EthernetMIIStatus"
    }
  },
  "EthernetMIIStatus": [
    {
      "InterfaceIndex": 4,
      "Name": "mgt0",
      "MauIndex": 1,
      "Control": "0x3140",
      "Status": "0x796D",
      "PhyIdentifierHigh": "0x0141",
      "PhyIdentifierLow": "0x0CB1",
      "AutoNegAdvertisement": "0x0DE1",
      "AutoNegLinkPartnerAbility": "0xC5E1",
      "AutoNegExpansion": "0x000D",
      "AutoNegNextPageTransmit": "",
      "AutoNegLinkPartnerRecNextPage": "",
      "MasterSlaveControlRegister": "0xFE1",
      "MasterSlaveStatusRegister": "0x3800",
      "PSEControl": "",
      "PSEStatus": "",
      "MMDAccessControl": "",
      "MMDAccessAddressData": "",
      "ExtendedStatus": "0x3000",
      "Vendor16": "",
      "Vendor17": "",
      "Vendor18": "",
      "Vendor19": "",
      "Vendor20": "",
      "Vendor21": "",
      "Vendor22": "",
      "Vendor23": "",
      "Vendor24": "",
      "Vendor25": "",
      "Vendor26": "",
      "Vendor27": "",
      "Vendor28": "",
      "Vendor29": "",
      "Vendor30": "",
      "Vendor31": ""
    },
    .....
    {
      "InterfaceIndex": 15,
      "Name": "eth17",
      "MauIndex": 1,
      "Control": "0x1840",
      "Status": "0x7949",
      "PhyIdentifierHigh": "0x0154",
      "PhyIdentifierLow": "0x03B1",
      "AutoNegAdvertisement": "0x0DE1",
      "AutoNegLinkPartnerAbility": "0x0000",
      "AutoNegExpansion": "0x0004",
      "AutoNegNextPageTransmit": "0x2001",
      "AutoNegLinkPartnerRecNextPage": "0x0000",
      "MasterSlaveControlRegister": "0x0600",
      "MasterSlaveStatusRegister": "0x0000",
      "PSEControl": "0x0000",
      "PSEStatus": "0x0000",
      "MMDAccessControl": "0x0000",
      "MMDAccessAddressData": "0x0000",
      "ExtendedStatus": "0x3000",
      "Vendor16": "0x0000",
      "Vendor17": "0x000A",
    }
  ]
}
```

```

    "Vendor18": "0x840A",
    "Vendor19": "0x1041",
    "Vendor20": "0x0000",
    "Vendor21": "0x000C",
    "Vendor22": "0xFF08",
    "Vendor23": "0x3048",
    "Vendor24": "0x0000",
    "Vendor25": "0x0000",
    "Vendor26": "0x0300",
    "Vendor27": "0x0003",
    "Vendor28": "0x210A",
    "Vendor29": "0x1F55",
    "Vendor30": "0x0000",
    "Vendor31": "0x03FC"
  }
]
}

```

[/mgmt/status/default/FailureNotificationStatus2](#)

You can use the HTTP GET method with the `/mgmt/status/default/FailureNotificationStatus2` resource to request a history of recently generated error reports.

GET

Use the HTTP GET method with the `/mgmt/status/default/FailureNotificationStatus2` resource to request information about generated failure notifications. A report is only available when the upload report option is enabled.

- [Resource URL](#)
- [“Request headers” on page 1050](#)
- [Request body format](#)
- [“Security requirements” on page 1050](#)
- [Response status codes](#)
- [“Response headers” on page 1051](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/FailureNotificationStatus2
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response lists recently generated reports:

Date

String.

Date and time the error report was created, in the format of `YYYYMMDDHHMMSSMILTMZ` where MIL is millisecond and TMZ is timezone.

Reason

String.

Event that triggered the creation of this error report:

- **crash** System crash occurred
- **watchdog** System deadlock situation occurred
- **low-memory** System crash occurred due to memory exhaustion
- **low-descriptor** System crash or deadlock occurred due to file descriptor exhaustion
- **user-requested** User requested must-gather via generate error report
- **on-startup** User requested on startup, reboot and reload using the failure notification always-on-startup configuration toggle
- **on-shutdown** User requested on shutdown using the failure notification always-on-shutdown configuration toggle
- **throttle** System throttler initiated reload
- **descriptor-error** System crash occurred due to file descriptor used or closed after already closed
- **unknown** System generated the report

UploadStatus

String.

Upload attempted success or failure status.

- **success** Error report creation and upload successful
- **failure** Error report upload or creation failed
- **inProgress** Error report upload is in-progress and will complete shortly

Location

String.

Location the error report was written to.

Examples

The following example requests information about recently uploaded reports.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/FailureNotificationStatus2
```

The following listing shows a received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/FailureNotificationStatus2"
    },
    "doc": {
      "href": "/mgmt/docs/status/FailureNotificationStatus2"
    }
  },
  "FailureNotificationStatus2": [
    {
      "Date": "20160507223218839BST",
      "Reason": "user-requested",
      "UploadStatus": "failure",
      "Location": "temporary:///error-report.7800537.20160507223218839BST.txt.gz"
    },
    {
      "Date": "20160809125843057BST",
      "Reason": "user-requested",
      "UploadStatus": "in-progress",
      "Location": "temporary:///error-report.7800537.20160809125843057BST.txt.gz"
    },
    {
      "Date": "20160809170210103BST",
      "Reason": "user-requested",
      "UploadStatus": "in-progress",
      "Location": "temporary:///error-report.7800537.20160809170210103BST.txt.gz"
    },
    {
      "Date": "20160809170717028BST",
      "Reason": "crash",
      "UploadStatus": "in-progress",
      "Location": "temporary:///error-report.7800537.20160809170717028BST.txt.gz"
    }
  ]
}
```

[/mgmt/status/default/FibreChannelHBAStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/FibreChannelHBAStatus` resource to request the status of the fibre channel host bus adapter.

Note: This request is available only on M2000 and M2001 appliances.

GET

Use the HTTP GET method with the `/mgmt/status/default/FibreChannelHBAStatus` resource to request the status of the fibre channel host bus adapter.

Note: This request is available only on M2000 and M2001 appliances.

- [Resource URL](#)
- [“Request headers” on page 1053](#)
- [Request body format](#)
- [“Security requirements” on page 1053](#)
- [Response status codes](#)
- [“Response headers” on page 1053](#)

- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/FibreChannelHBAStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated”](#) on page 423 and [“User authorization, credential mapping, and access profiles”](#) on page 423.

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes status information from the fibre channel host bus adapters. Each object contains the following fields:

HBAInstance

String.

Host bus adapter instance.

OpState

String.

Configuration object operational state. Can indicate up or down.

WorldWidePortNumber

String.

The World Wide Port Number for this HBA instance.

PortState

String.

The state of the port for this HBA instance.

PortSpeed

String.

The speed of the port for this HBA instance.

PortType

String.

The type of the port for this HBA instance.

SupportedSpeeds

The speeds supported by the port for this HBA instance.

Examples

The following example requests the status of fibre channel host bus adapters on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/FibreChannelHBAStatus
```

The following listing shows sections from the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/FibreChannelHBAStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/FibreChannelHBAStatus"
    }
  },
  "FibreChannelHBAStatus": [
    {
      "HBAInstance": "fch1",
      "OpState": "down",
      "WorldWidePortNumber": "",
      "PortState": "",
      "PortSpeed": "",
      "PortType": "",
      "SupportedSpeeds": ""
    },
    {
      "HBAInstance": "fch2",
      "OpState": "down",
      "WorldWidePortNumber": "",
      "PortState": "",
      "PortSpeed": "",
      "PortType": "",
      "SupportedSpeeds": ""
    }
  ]
}
```

[/mgmt/status/default/FibreChannelLuns](#)

You can use the HTTP GET method with the `/mgmt/status/default/FibreChannelLuns` resource to request a list of discovered fibre channel LUNs.

Note: This request is available only on M2000 and M2001 appliances.

GET

Use the HTTP GET method with the `/mgmt/status/default/FibreChannelLuns` resource to request a list of discovered fibre channel LUNs.

Note: This request is available only on M2000 and M2001 appliances.

- [Resource URL](#)
- [“Request headers” on page 1055](#)
- [Request body format](#)
- [“Security requirements” on page 1055](#)
- [Response status codes](#)
- [“Response headers” on page 1056](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/FibreChannelLuns
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about fibre channel LUNs. Each object contains the following fields:

luid

String.

Unique ID of a SAN LUN.

hba

String.

Host Bus Adapter instance.

lun

String.

Logical unit number assigned by the operating system to the logical unit.

wwpn

String.

World-wide name of the target port to be queried for information about the logical unit.

Examples

The following example requests a list of fibre channel LUNs on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/FibreChannelLuns
```

The following listing shows sections from the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/FibreChannelLuns"
    },
    "doc": {
      "href": "/mgmt/docs/status/FibreChannelLuns"
    }
  },
  "FibreChannelLuns": [
    {
      "luid": "600507680181804D9800000000001B48",
      "hba": "fch2",
      "lun": 5,
      "wwpn": "50:05:07:68:01:10:26:65"
    },
    {
      "luid": "600507680181804D9800000000001B48",
      "hba": "fch2",
      "lun": 5,
      "wwpn": "50:05:07:68:01:10:27:05"
    },
    {
      "luid": "600507680181804D9800000000001B48",
      "hba": "fch1",
      "lun": 5,
      "wwpn": "50:05:07:68:01:30:26:65"
    },
    {
      "luid": "600507680181804D9800000000001B48",
      "hba": "fch1",
      "lun": 5,
      "wwpn": "50:05:07:68:01:30:27:05"
    }
  ]
}
```

```
]
}
```

[/mgmt/status/default/FibreChannelVolumeStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/FibreChannelVolumeStatus` resource to request the status of Fibre Channel volumes.

Note: This request is available only on M2000 and M2001 appliances.

GET

Use the HTTP GET method with the `/mgmt/status/default/FibreChannelVolumeStatus` resource to request the status of Fibre Channel volumes.

Note: This request is available only on M2000 and M2001 appliances.

- [Resource URL](#)
- [“Request headers” on page 1057](#)
- [Request body format](#)
- [“Security requirements” on page 1057](#)
- [Response status codes](#)
- [“Response headers” on page 1058](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/FibreChannelVolumeStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes status information from fibre channel volumes. Each object contains the following fields:

volume

String.

Name of the Fibre Channel volume.

luid

String.

Unique LUN ID used by the volume.

OpStatus

String.

Operational state of the volume object.

linkStatus

String.

Status of the link to the Fibre Channel volume.

Examples

The following example requests the status of Fibre Channel volumes on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/FibreChannelVolumeStatus
```

The following listing shows sections from the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/FibreChannelVolumeStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/FibreChannelVolumeStatus"
    }
  },
  "FibreChannelVolumeStatus": {
    "volume": "FCVolume1",
    "luid": "600507680181804D9800000000001B48",
    "opState": "up",
    "linkStatus": "up"
  }
}
```

[/mgmt/status/default/FilesystemStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/FilesystemStatus` resource to request the current usage of storage capacity on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/FilesystemStatus` resource to request the current usage of storage capacity on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1059](#)
- [Request body format](#)
- [“Security requirements” on page 1059](#)
- [Response status codes](#)
- [“Response headers” on page 1059](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/FilesystemStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes usage information for storage on the appliance. Each object contains the following fields:

FreeEncrypted

Integer.

Free, or unused and available, encrypted storage space on the appliance. All files placed in encrypted storage are automatically encrypted for security by the appliance.

TotalEncrypted

Integer.

Total encrypted storage space on the appliance (the maximum capacity). All files placed in encrypted storage are automatically encrypted for security by the appliance.

FreeTemporary

Integer.

Free, or unused and available, temporary storage space on the appliance. The appliance uses temporary storage for processing, logging and debugging.

TotalTemporary

Integer.

Total temporary storage space on the appliance. The appliance uses temporary storage for processing, logging and debugging.

FreeInternal

Integer.

Free, or unused and available, internal storage space on the appliance. The appliance uses internal storage for import, export, firmware upgrades and debug captures.

TotalInternal

Integer.

Total internal storage space on the appliance. The appliance uses internal storage for import, export, firmware upgrades and debug captures.

Examples

The following example requests usage information for storage on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/FileSystemStatus
```

The following listing shows sections from the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/FileSystemStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/FileSystemStatus"
    }
  },
  "FileSystemStatus": {
    "FreeEncrypted": 8683,
    "TotalEncrypted": 15008,
    "FreeTemporary": 4051,
    "TotalTemporary": 4096,
    "FreeInternal": 3071,
    "TotalInternal": 3072
  }
}
```

[/mgmt/status/default/FirmwareStatus2](#)

You can use the HTTP GET method with the `/mgmt/status/default/FirmwareStatus2` resource to request information about the current firmware installation.

GET

Use the HTTP GET method with the `/mgmt/status/default/FirmwareStatus2` resource to request information about the current firmware installation.

- [Resource URL](#)
- [“Request headers” on page 1061](#)
- [Request body format](#)
- [“Security requirements” on page 1061](#)
- [Response status codes](#)
- [“Response headers” on page 1062](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/FirmwareStatus2
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the current firmware installation. Each object contains the following fields:

Type

String.

The type of the firmware installation type. Can be primary or secondary. The secondary firmware image is used when rolling back the firmware.

- **Primary** Primary firmware install
- **Secondary** Secondary firmware install

Version

String.

The version of this firmware image

Build

String.

The specific build of this firmware image.

BuildDate

String.

Date of the firmware build.

DeliveryType

String.

Indicates the delivery type.

- **LTS** Long-term support (LTS) releases are intended for customers who demand the highest level of stability. Updates during the support period contain fixes to identified problems and vulnerabilities.
- **CD** Continuous delivery (CD) releases add new function on a regular cadence and are intended for customers who want to exploit the latest features and capabilities without waiting for the next LTS release. Each CD release contains fixes to identified problems and vulnerabilities.

InstallDate

String.

Date that this firmware was installed.

BootCount

Integer.

The number of times the firmware image was restarted through an appliance reboot or a firmware reload. The count is from the initial firmware load on the appliance till the current time. The count is independent of firmware version.

Examples

The following example requests information about the appliance firmware.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/FirmwareStatus2
```

The following listing shows sections from the received response:

```
{  
  "_links": {  
    "self": {
```

```

    "href": "/mgmt/status/default/FirmwareStatus2"
  },
  "doc": {
    "href": "/mgmt/docs/status/FirmwareStatus2"
  }
},
"FirmwareStatus2": {
  "Type": "primary",
  "Version": "MQ00.9.1.2.0",
  "Build": "306735mq",
  "BuildDate": "2019/02/27 12:09:32",
  "DeliveryType": "CD",
  "InstallDate": "2019/03/07 14:13:15",
  "BootCount": 6
}
}
}

```

[/mgmt/status/default/FirmwareStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/FirmwareStatus` resource to request information about the current firmware installation.

GET

Use the HTTP GET method with the `/mgmt/status/default/FirmwareStatus2` resource to request information about the current firmware installation.

- [Resource URL](#)
- [“Request headers” on page 1063](#)
- [Request body format](#)
- [“Security requirements” on page 1063](#)
- [Response status codes](#)
- [“Response headers” on page 1064](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/FirmwareStatus2
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers**Content-Type**

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the current firmware installation. Each object contains the following fields:

Type

String.

The type of the firmware installation type. Can be primary or secondary. The secondary firmware image is used when rolling back the firmware.

- **Primary** Primary firmware install
- **Secondary** Secondary firmware install

Version

String.

The version of this firmware image

Build

String.

The specific build of this firmware image.

BuildDate

String.

Date of the firmware build.

InstallDate

String.

Date that this firmware was installed.

BootCount

Integer.

The number of times the firmware image was restarted through an appliance reboot or a firmware reload. The count is from the initial firmware load on the appliance till the current time. The count is independent of firmware version.

Examples

The following example requests information about the appliance firmware.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/FirmwareStatus2
```

The following listing shows sections from the received response:

```

{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/FirmwareStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/FirmwareStatus"
    }
  },
  "FirmwareStatus": {
    "Type": "primary",
    "Version": "MQ00.9.1.1.0",
    "Build": "mq-warrior.304452",
    "BuildDate": "2018/11/12 11:38:40",
    "InstallDate": "2018/11/13 09:40:28",
    "BootCount": 427
  }
}

```

[/mgmt/status/default/FirmwareVersion3](#)

You can use the HTTP GET method with the `/mgmt/status/default/FirmwareVersion3` resource to list the version of the firmware and its components as well as the machine and model type for the appliance. You need this information to report problems.

GET

Use the HTTP GET method with the `/mgmt/status/default/FirmwareVersion3` resource to list the version of the firmware and its components as well as the machine and model type for the appliance.

- [Resource URL](#)
- [“Request headers” on page 1065](#)
- [Request body format](#)
- [“Security requirements” on page 1065](#)
- [Response status codes](#)
- [“Response headers” on page 1066](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/FirmwareVersion3
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the current firmware version. Each object contains the following fields:

Serial

String.

The serial number of the appliance. This number is unique to each appliance.

Version

String.

The version of this firmware image

Level

String.

The level of the firmware.

Build

String.

The specific build of this firmware image.

BuildDate

String.

Date of the firmware build.

DeliveryType

String.

Indicates the delivery type.

- **LTS** Long-term support (LTS) releases are intended for customers who demand the highest level of stability. Updates during the support period contain fixes to identified problems and vulnerabilities.
- **CD** Continuous delivery (CD) releases add new function on a regular cadence and are intended for customers who want to exploit the latest features and capabilities without waiting for the next LTS release. Each CD release contains fixes to identified problems and vulnerabilities.

WatchdogBuild

String.

The build number of the system health monitor.

InstalledDPOS

String.

The version number of the appliance manager.

RunningDPOS

String.

The version number of the appliance manager.

XMLAccelerator

String.

MachineType

String.

The IBM machine type of the appliance.

ModelType

String.

The IBM model type of the appliance.

Tenant

String.

not used

Examples

The following example requests information about the appliance firmware.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/FirmwareVersion3
```

The following listing shows sections from the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/FirmwareVersion3"
    },
    "doc": {
      "href": "/mgmt/docs/status/FirmwareVersion3"
    }
  },
  "FirmwareVersion3": {
    "Serial": "1300014",
    "Version": "MQ00.9.1.2.0",
    "Level": "p912-L190305",
    "Build": "306735mq",
    "BuildDate": "2019/02/27 12:09:32",
    "DeliveryType": "CD",
    "WatchdogBuild": "MQ00.9.1.2.0",
    "InstalledDPOS": "MQ00.9.1.2.0",
    "RunningDPOS": "MQ00.9.1.2.0",
    "XMLAccelerator": "embedded",
    "MachineType": "8441",
    "ModelType": "54X",
    "Tenant": ""
  }
}
```

[/mgmt/status/default/IGMPStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/IGMPStatus` resource to request status information about Internet Group Management Protocol (IGMP) on interfaces on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/IGMPStatus` resource to request status information about Internet Group Management Protocol (IGMP) on interfaces on the appliance.

- [Resource URL](#)

- [“Request headers” on page 1068](#)
- [Request body format](#)
- [“Security requirements” on page 1068](#)
- [Response status codes](#)
- [“Response headers” on page 1068](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/IGMPStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about IGMP status of appliance interfaces. Each object contains the following fields:

InterfaceIdx

Integer.

The index of the local interface associated with the group.

Interface

String.

The local interface associated with the group.

InterfaceType

String.

The type of device interface used for this route. Ethernet indicates that it is on the base Ethernet MAC interface. VLAN indicates that it is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.

- **Ethernet** Indicates that it is on the base Ethernet MAC interface.
- **Tunnel** Indicates that it is on a encapsulation interface.
- **VLAN** Indicates that it is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.
- **Aggregate** Indicates that it is on an IEEE 802.3ad Link Aggregation composed of one of the base Ethernet MAC interface.
- **Other** 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.

RefCount

Integer.

Reference count.

Group

String.

IGMP group.

Examples

The following example requests information about the IGMP status of appliance interfaces.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/IGMPStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/IGMPStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/IGMPStatus"
    }
  },
  "IGMPStatus": [
    {
      "InterfaceIdx": 1,
      "Interface": "loopback",
      "InterfaceType": "Ethernet",
      "RefCount": 1,
      "Group": "224.0.0.1"
    },
    {
      "InterfaceIdx": 1,
      "Interface": "loopback",
      "InterfaceType": "Ethernet",
      "RefCount": 1,
      "Group": "239.1.1.1"
    },
    {
      "InterfaceIdx": 1,
      "Interface": "loopback",

```

```

    "InterfaceType": "Ethernet",
    "RefCount": 1,
    "Group": "239.1.1.2"
  },
  {
    "InterfaceIdx": 4,
    "Interface": "mgt0",
    "InterfaceType": "Ethernet",
    "RefCount": 1,
    "Group": "224.0.0.1"
  },
  {
    "InterfaceIdx": 6,
    "Interface": "eth20",
    "InterfaceType": "Ethernet",
    "RefCount": 1,
    "Group": "224.0.0.1"
  },
  {
    "InterfaceIdx": 10,
    "Interface": "eth21",
    "InterfaceType": "Ethernet",
    "RefCount": 1,
    "Group": "224.0.0.1"
  },
  {
    "InterfaceIdx": 11,
    "Interface": "eth13",
    "InterfaceType": "Ethernet",
    "RefCount": 1,
    "Group": "224.0.0.1"
  }
]
}

```

[/mgmt/status/default/IPAddressStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/IPAddressStatus` resource to request status information about interfaces on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/IPAddressStatus` resource to request status information about interfaces on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1070](#)
- [Request body format](#)
- [“Security requirements” on page 1071](#)
- [Response status codes](#)
- [“Response headers” on page 1071](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/IGMPStatus
```

Where `port` is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about appliance interfaces. Each object contains the following fields:

Name

String.

The name of the interface.

Index

Integer.

The value of the SNMP `ifIndex` column in the `ifTable` for the interface.

IPVersion

String.

Indicates the version of IP Address for the interface.

- **unknown** An unknown address type.
- **ipv4** An IPv4 address.
- **ipv6** An IPv6 address.
- **ipv4z** A nonglobal IPv4 address with a zone index.
- **ipv6z** A nonglobal IPv6 address with a zone index.
- **dns** A DNS domain name

PrefixLength

Integer.

Indicates the prefix length for the interface. The prefix length represents the configured netmask for the interface.

IP

String.

The IP address for the interface.

Examples

The following example requests information about the appliance interfaces.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/IPAddressStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/IPAddressStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/IPAddressStatus"
    }
  },
  "IPAddressStatus": [
    {
      "Name": "lo",
      "Index": 1,
      "IPVersion": "ipv4",
      "PrefixLength": 8,
      "IP": "127.0.0.1"
    },
    {
      "Name": "lo",
      "Index": 1,
      "IPVersion": "ipv6",
      "PrefixLength": 128,
      "IP": "::1"
    },
    {
      "Name": "mgt0",
      "Index": 4,
      "IPVersion": "ipv4",
      "PrefixLength": 24,
      "IP": "192.0.2.136"
    },
    {
      "Name": "eth20",
      "Index": 6,
      "IPVersion": "ipv4",
      "PrefixLength": 16,
      "IP": "172.21.136.0"
    },
    {
      "Name": "eth21",
      "Index": 10,
      "IPVersion": "ipv4",
      "PrefixLength": 24,
      "IP": "192.168.253.136"
    },
    {
      "Name": "eth13",
      "Index": 11,
      "IPVersion": "ipv4",
      "PrefixLength": 24,
      "IP": "192.168.251.136"
    }
  ]
}
```

[/mgmt/status/default/IPMISelEvents](#)

You can use the HTTP GET method with the `/mgmt/status/default/IPMISelEvents` resource to request a list of the Intelligent Platform Management Interface (IPMI) System Event Log (SEL) events that the Baseboard Management Controller (BMC) collected.

GET

Use the HTTP GET method with the `/mgmt/status/default/IPMISelEvents` resource to request a list of the Intelligent Platform Management Interface (IPMI) System Event Log (SEL) events that the Baseboard Management Controller (BMC) collected.

- [Resource URL](#)
- [“Request headers” on page 1073](#)
- [Request body format](#)
- [“Security requirements” on page 1073](#)
- [Response status codes](#)
- [“Response headers” on page 1074](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/IPMISelEvents
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes IPMI System Event Log (SEL) events. Each object contains the following fields:

Index

Integer.

The simple sequential index into this table.

Timestamp

String.

Timestamp of when this event was written to the SEL in the IPMI BMC.

RecordType

String.

The record type in hex of this record.

SensorType

String.

The type of sensor that generated this event.

- reserved
- temperature
- voltage
- current
- fan
- physicalSecurity
- platformSecurityViolationAttempt
- processor
- powerSupply
- powerUnit
- coolingDevice
- otherUnitsBasedSensor
- memory
- driveSlot
- postMemoryResize
- systemFirmwareProgress
- eventLoggingDisabled
- watchdog1
- systemEvent
- criticalInterrupt
- button
- moduleOrBoard
- microcontrollerOrCoprocesor
- addInCard
- chassis
- chipSet

- otherFru
- cableOrInterconnect
- terminator
- systemBootRestartInitiated
- bootError
- osBoot
- osCriticalStop
- slotOrConnector
- systemAcpiPowerStatus
- watchdog2
- platformAlert
- entityPresence
- monitorAsicIc
- lan
- managmentSubsystemHealth
- battery
- sessionAudit
- versionChange
- fruState
- kernel

SensorNumber

String.

The internal IPMI BMC sensor number of the sensor that generated this event.

SensorName

String.

The name of the sensor that generated this event.

EventReadingTypeCode

String.

The type of event transition that the sensor made that caused this event record to be added to the BMC SEL. (Not all of the listed values are used by the appliance.)

- None
- undefined
- lowerNonCriticalGoingLow
- lowerNonCriticalGoingHigh
- lowerCriticalGoingLow
- lowerCriticalGoingHigh
- lowerNonRecoverableGoingLow
- lowerNonRecoverableGoingHigh
- upperNonCriticalGoingLow
- upperNonCriticalGoingHigh
- upperCriticalGoingLow
- upperCriticalGoingHigh
- upperNonRecoverableGoingLow
- upperNonRecoverableGoingHigh

- transitionToIdle
- transitionToActive
- transitionToBusy
- stateDeasserted
- stateAsserted
- predictiveFailureDeasserted
- predictiveFailureAsserted
- limitNotExceeded
- limitExceeded
- performanceMet
- performanceLags
- transitionToOK
- transitionToNonCriticalFromOK
- transitionToCriticalFromLessSevere
- transitionToNonRecoverableFromLessSevere
- transitionToNonCriticalFromMoreSevere
- transitionToCriticalFromNonRecoverable
- transitionToNonRecoverable
- monitor
- informational
- deviceRemovedOrAbsent
- deviceInsertedOrPresent
- deviceDisabled
- deviceEnabled
- transitionToRunning
- transitionToInTest
- transitionToPowerOff
- transitionToOnLine
- transitionToOffLine
- transitionToOffDuty
- transitionToDegraded
- transitionToPowerSave
- installError
- fullyRedundant
- redundancyLost
- redundancyDegraded
- nonRedundantSufficientResourcesFromRedundant
- nonRedundantSufficientResourcesFromInsufficient
- nonRedundantInsufficientResources
- redundancyDegradedFromFullyRedundant
- redundancyDegradedFromNonRedundant
- acpiD0PowerState
- acpiD1PowerState

- acpiD2PowerState
- acpiD3PowerState
- temperature
- voltage
- current
- fan
- generalChassisIntrusion
- driveBayIntrusion
- ioCardAreaIntrusion
- processorAreaIntrusion
- lanLeashLost
- unauthorizedDockUndock
- fanAreaIntrusion
- secureModeViolationAttempt
- preBootPasswordViolationUser
- preBootPasswordViolationSetup
- preBootPasswordViolationNetworkBoot
- otherPreBootPasswordViolation
- outOfBandAccessPasswordViolation
- ierr
- thermalTrip
- frb1BistFailure
- frb2HangInPostFailure
- frb3ProcessorStartupInitializationFailure
- configurationError
- smbiosUncorrectableCpuComplexError
- processorPresenceDetected
- processorDisabled
- terminatorPresenceDetected
- processorAutomaticallyThrottled
- powerSupplyPresenceDetected
- powerSupplyFailureDetected
- powerSupplyPredictiveFailure
- powerSupplyAcLost
- powerSupplyAcLostOrOutOfRange
- powerSupplyAcOutOfRangeButPresent
- powerSupplyConfigurationError
- powerUnitPowerOff
- powerUnitPowerCycle
- powerUnit240VaPowerDown
- powerUnitInterlockPowerDown
- powerUnitAcLost
- powerUnitSoftPowerControlFailure

- powerUnitFailureDetected
- powerUnitPredictiveFailure
- correctableEcc
- uncorrectableEcc
- parity
- memoryScrubFailed
- memoryDeviceDisabled
- correctableEccLoggingLimitReached
- memoryPresenceDetected
- memoryConfigurationError
- memorySpare
- memoryAutomaticallyThrottled
- memoryCriticalOvertemperature
- drivePresence
- driveFault
- drivePredictiveFailure
- driveHotSpare
- driveConsistencyCheckInProgress
- driveInCriticalArray
- driveInFailedArray
- driveRebuildInProgress
- driveRebuildAborted
- systemFirmwareError
- systemFirmwareHang
- systemFirmwareProgress
- correctableMemoryErrorLoggingDisabled
- eventLoggingDisabled
- logAreaReset
- allEventLoggingDisabled
- selFull
- selAlmostFull
- biosWatchdogReset
- osWatchdogReset
- osWatchdogShutDown
- osWatchdogPowerDown
- osWatchdogPowerCycle
- osWatchdogNmiDiagnosticInterrupt
- osWatchdogExpiredStatusOnly
- osWatchdogPreTimeoutInterrupt
- systemReconfigured
- oemSystemBootEvent
- undeterminedSystemHardwareFailure
- entryAddedToAuxiliaryLog

- pefAction
- timestampClockSync
- frontPanelNmiInterrupt
- busTimeout
- ioChannelCheckNmi
- softwareNmi
- pciPerr
- pciSerr
- eisaFailSafeTimeout
- busCorrectableError
- busUncorrectableError
- fatalNmi
- busFatalError
- busDegraded
- powerButtonPressed
- sleepButtonPressed
- resetButtonPressed
- fruLatchOpen
- fruServiceRequestButton
- softPowerControlFailure
- initiatedByPowerUp
- initiatedByHardReset
- initiatedByWarmReset
- userRequestedPxeBoot
- automaticBootToDiagnostic
- noBootableMedia
- nonBootableDiskette
- pxeServerNotFound
- invalidBootSector
- timeoutWaitingForUserSelectionOfBootSource
- aBootCompleted
- cBootCompleted
- pxeBootCompleted
- diagnosticBootCompleted
- cdromBootCompleted
- romBootCompleted
- bootCompleted
- stopDuringOsLoadInitialization
- runTimeStop
- slotFaultStatusAsserted
- slotIdentifyStatusAsserted
- slotDeviceInstalled
- slotReadyForDeviceInstallation

- slotReadyForDeviceRemoval
- slotPowerIsOff
- slotDeviceRemovalRequest
- slotInterlockAsserted
- slotIsDisabled
- slotHoldsSpare
- acpiS0G0PowerState
- acpiS1PowerState
- acpiS2PowerState
- acpiS3PowerState
- acpiS4PowerState
- acpiS5G2PowerState
- acpiS4S5SoftOff
- acpiG3PowerState
- acpiSleepingS1S2S3
- acpiG1PowerState
- acpiS5OverrideState
- acpiLegacyOnState
- acpiLegacyOffState
- acpiUnknown
- watchdog2TimeExpired
- watchdog2HardReset
- watchdog2PowerDown
- watchdog2PowerCycle
- watchdog2TimerInterrupt
- platformGeneratedPage
- platformGeneratedLanAlert
- platformGeneratedEventTrap
- platformGeneratedSnmpTrap
- entityPresent
- entityAbsent
- entityDisabled
- lanHeartbeatLost
- lanHeartbeat
- sensorAccessDegradedOrUnavailable
- controllerAccessDegradedOrUnavailable
- managementControllerOffLine
- managementControllerUnavailable
- batteryLow
- batteryFailed
- batteryPresenceDetected
- sessionActivated
- sessionDeactivated

- hardwareVersionChange
- firmwareVersionChange
- hardwareIncompatibility
- firmwareIncompatibility
- entityIsUnsupportedHardwareVersion
- entityContainsUnsupportedFirmwareVersion
- hardwareChangeSuccessful
- firmwareChangeSuccessful
- fruNotInstalled
- fruInactive
- fruActivationRequested
- fruActivationInProgress
- fruActive
- fruDeactivationRequested
- fruDeactivationInProgress
- fruCommunicationLost
- kernelPanic

EventData2

String.

Additional data for some event types from the Event Data 2 field of the SEL entry. (Not all of the listed values are used by the appliance.)

- none
- unspecified
- noSystemMemoryInstalled
- noUsableSystemMemory
- unrecoverableHardDiskDeviceFailure
- unrecoverableSystemBoardFailure
- unrecoverableDisketteSubsystemFailure
- unrecoverableHardDiskControllerFailure
- unrecoverableKeyboardFailure
- removableBootMediaNotFound
- unrecoverableVideoControllerFailure
- noVideoDeviceDetected
- firmwareRomCorruptionDetected
- cpuVoltageMismatch
- cpuSpeedMatchingFailure
- memoryInitialization
- hardDiskInitialization
- secondaryProcessorInitialization
- userAuthentication
- userInitiatedSystemStartup
- usbResourceConfiguration
- pciResourceConfiguration
- optionRomInitialization

- videoInitialization
- cacheInitialization
- smbusInitialization
- keyboardControllerInitialization
- embeddedManagementControllerInitialization
- dockingStationAttachment
- enablingDockingStation
- dockingStationEjection
- disablingDockingStation
- callingOperatingSystemWakeupVector
- startingOperatingSystemBoot
- baseboardInitialization
- reserved
- floppyInitialization
- keyboardTest
- pointingDeviceTest
- primaryProcessorInitialization
- interruptTypeSMI
- interruptTypeNMI
- interruptTypeMessagingInterrupt

EventDirection

String.

The direction of the sensor event that caused this event record to be added to the BMC SEL.

- none
- asserted
- deasserted

Extra

String.

Extra information used for some event types.

Examples

The following example requests information about IPMI System Event Log (SEL) events that have occurred.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/IPMISelEvents
```

The following listing shows sections from an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/IPMISelEvents"
    },
    "doc": {
      "href": "/mgmt/docs/status/IPMISelEvents"
    }
  },
  "IPMISelEvents": [
    {
      "Index": 1,
```

```

    "Timestamp": "2016/04/11 13:48:58",
    "RecordType": "0x0002",
    "SensorType": "systemEvent",
    "SensorNumber": "0x0000",
    "SensorName": "",
    "EventReadingTypeCode": "timestampClockSync",
    "EventData2": "none",
    "EventDirection": "asserted",
    "Extra": ""
  },
  {
    "Index": 2,
    "Timestamp": "2016/04/11 13:48:58",
    "RecordType": "0x0002",
    "SensorType": "systemEvent",
    "SensorNumber": "0x0083",
    "SensorName": "SystemEvent",
    "EventReadingTypeCode": "timestampClockSync",
    "EventData2": "none",
    "EventDirection": "asserted",
    "Extra": ""
  },
  {
    "Index": 3,
    "Timestamp": "2016/04/11 13:49:06",
    "RecordType": "0x0002",
    "SensorType": "systemEvent",
    "SensorNumber": "0x0083",
    "SensorName": "SystemEvent",
    "EventReadingTypeCode": "timestampClockSync",
    "EventData2": "none",
    "EventDirection": "asserted",
    "Extra": ""
  },
  ...
  {
    "Index": 154,
    "Timestamp": "2018/11/06 10:34:40",
    "RecordType": "0x0002",
    "SensorType": "systemEvent",
    "SensorNumber": "0x0083",
    "SensorName": "SystemEvent",
    "EventReadingTypeCode": "timestampClockSync",
    "EventData2": "none",
    "EventDirection": "asserted",
    "Extra": ""
  },
  {
    "Index": 155,
    "Timestamp": "2018/11/06 10:34:40",
    "RecordType": "0x0002",
    "SensorType": "systemEvent",
    "SensorNumber": "0x0000",
    "SensorName": "",
    "EventReadingTypeCode": "timestampClockSync",
    "EventData2": "none",
    "EventDirection": "asserted",
    "Extra": ""
  }
]
}

```

[/mgmt/status/default/LicenseStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/LicenseStatus` resource to request information about available and enabled features.

GET

Use the HTTP GET method with the `/mgmt/status/default/LicenseStatus` resource to request information about available and enabled features.

- [Resource URL](#)
- [“Request headers” on page 1084](#)
- [Request body format](#)
- [“Security requirements” on page 1084](#)

- [Response status codes](#)
- [“Response headers” on page 1084](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/LicenseStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about appliance features. Each object contains the following fields:

Feature

String.

The name of the feature.

- **DataGlue** Binary processing support

- **RaidVolume** Hard disk array
- **LocateLED** Locate LED light available
- **FibreChannel** Fibre Channel (M2000 and M2001 only)
- **IPMI** Intelligent Platform Management Interface
- **RaidVolumeSr** ServeRAID RAID support
- **IntrusionDetection** Intrusion Detection support
- **IPMI-LAN** Intelligent Platform Management Interface LAN channel
- **IPv6** IPv6 support
- **RealHardware** System running on real hardware
- **RealRaid** Hard Disk Array backed by a real hardware
- **Language** Language support
- **IMQA** Base IBM MQ Appliance

Enabled

The enabled status of this feature. Set to **Yes**, **No**, or **Error**.

Available

The availability of this feature. A feature is unavailable when the installed firmware image does not include the code to support the feature. Set to **Yes** or **No**.

Examples

The following example requests information about appliance features.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/LicenseStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/LicenseStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/LicenseStatus"
    }
  },
  "LicenseStatus": [
    {
      "Feature": "DataGlue",
      "Enabled": "Yes",
      "Available": "Yes"
    },
    {
      "Feature": "RaidVolume",
      "Enabled": "Yes",
      "Available": "Yes"
    },
    {
      "Feature": "LocateLED",
      "Enabled": "Yes",
      "Available": "Yes"
    },
    {
      "Feature": "FibreChannel",
      "Enabled": "Yes",
      "Available": "Yes"
    },
    {
      "Feature": "IPMI",
      "Enabled": "Yes",
      "Available": "Yes"
    },
    {
      "Feature": "RaidVolumeSr",
```

```

    "Enabled": "Yes",
    "Available": "Yes"
  },
  {
    "Feature": "IntrusionDetection",
    "Enabled": "Yes",
    "Available": "Yes"
  },
  {
    "Feature": "IPMI-LAN",
    "Enabled": "Yes",
    "Available": "Yes"
  },
  {
    "Feature": "IPv6",
    "Enabled": "Yes",
    "Available": "Yes"
  },
  {
    "Feature": "RealHardware",
    "Enabled": "Yes",
    "Available": "Yes"
  },
  {
    "Feature": "RealRaid",
    "Enabled": "Yes",
    "Available": "Yes"
  },
  {
    "Feature": "Language",
    "Enabled": "Yes",
    "Available": "Yes"
  },
  {
    "Feature": "IMQA",
    "Enabled": "Yes",
    "Available": "Yes"
  }
]
}

```

[/mgmt/status/default/LinkAggregationStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/LinkAggregationStatus` resource to request statistics for aggregate interfaces.

GET

Use the HTTP GET method with the `/mgmt/status/default/LinkAggregationStatus` resource to request statistics for aggregate interfaces.

- [Resource URL](#)
- [“Request headers” on page 1086](#)
- [Request body format](#)
- [“Security requirements” on page 1087](#)
- [Response status codes](#)
- [“Response headers” on page 1087](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/LinkAggregationStatus
```

Where `port` is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about aggregated interfaces on the appliance. Each object contains the following fields:

Index

Integer.

Identifies the index that identifies the aggregate interface.

Name

String.

The name of the interface in the system kernel.

MgmtName

String.

Indicates the name of the interface in the configuration.

MTU

Integer.

Identifies the maximum transmission unit (MTU) for the aggregation.

Mode

String.

Indicates the mode for link aggregation.

- **ab** Active back up. Indicates that the mode is active-backup. One link is active, and the other link is standby.
- **LACP** LACP mode. Indicates that the mode uses the Link Aggregation Control Protocol (LACP). The IEEE 802.1AX-2008 standard defines LACP.
- **tlb** Transmit load balancing. Indicates that the mode uses transmit-based load balancing.
- **none** Indicates that the mode is not supported. Contact IBM Support.

Status

String.

Indicates the link state of the aggregate interface.

- **ok** The interface is responding to the network.
- **noLink** The interface is not responding to the network.

Primary

String.

Indicates which link, if any, is primary when the link aggregation mode is active-backup.

Active

Indicates the active link when the link aggregation mode is transmit-based load balancing or active-backup.

LACPHashPolicy

String.

With the LACP aggregation mode, indicates the distribution algorithm for outbound packets among the active physical interfaces. The policy should match the configuration.

LACPDURate

String.

With the LACP aggregation mode, indicates the transmission rate that members send LACP frames. The rate can be either fast or slow. In LACP, the frame uses the term LACPDU.

LACPSelectionPolicy

String.

With the LACP aggregation mode, indicates the selection policy. The policy should match the configuration.

LACPAggregatorID

String.

With the LACP aggregation mode, indicates the ID of the active aggregator.

LACPNumPorts

With the LACP aggregation mode, indicates the number of available ports for the active aggregator.

LACPActorKey

With the LACP aggregation mode, identifies the value of the operational key that the Actor assigned to the port.

LACPPartnerKey

String.

With the LACP aggregation mode, identifies the physical MAC address of the Partner.

LACPPartnerMAC

String.

The physical MAC address of the Partner.

MIIStatus

String.

Indicates the MII link state of the aggregation.

- **ok** The interface is responding to the network.
- **noLink** The interface is not responding to the network.

Examples

The following example requests information about aggregated interfaces on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/LinkAggregationStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/LinkAggregationStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/LinkAggregationStatus"
    }
  },
  "result": "No status retrieved."
}
```

[/mgmt/status/default/LinkAggregationMemberStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/LinkAggregationMemberStatus` resource to request a mapping of links to aggregates on the system.

GET

Use the HTTP GET method with the `/mgmt/status/default/LinkAggregationMemberStatus` resource to request a mapping of links to aggregates on the system.

- [Resource URL](#)
- [“Request headers” on page 1089](#)
- [Request body format](#)
- [“Security requirements” on page 1089](#)
- [Response status codes](#)
- [“Response headers” on page 1090](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/LinkAggregationMemberStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes mapping information about members of aggregated interfaces on the appliance. Each object contains the following fields:

AggregateIndex

Integer.

The index that identifies the link aggregation.

AggregateName

String.

The name of the interface in the system kernel.

MgmtName

String.

Indicates the name of the interface in the configuration.

MemberIndex

Integer.

The index that identifies the member interface.

MemberName

String.

The system name of the member interface.

AggregatorID

String.

Aggregator ID assigned to the interface.

Examples

The following example requests mapping information about members of aggregated interfaces on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/LinkAggregationMemberStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/LinkAggregationMemberStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/LinkAggregationMemberStatus"
    }
  },
  "result": "No status retrieved."
}
```

[/mgmt/status/default/LinkStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/LinkStatus` resource to request status for the interfaces on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/LinkStatus` resource to request status for the interfaces on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1091](#)
- [Request body format](#)
- [“Security requirements” on page 1091](#)
- [Response status codes](#)
- [“Response headers” on page 1092](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/LinkStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about interfaces on the appliance. Each object contains the following fields:

Name

String.

The name of the interface in the configuration.

Index

Integer.

The if index for the interface.

Status

String.

the current link state for the interface.

- **ok** The interface is responding to the network.
- **noLink** The interface is not responding to the network.

Mode

String.

Identifies the actual fixed, configured, or negotiated Ethernet PHY mode, including speed and duplex.

- **10BASE-T-HD** PHY is running 10BASE-T (10 mbps) in half-duplex mode.
- **10BASE-T-FD** PHY is running 10BASE-T (10 mbps) in full-duplex mode.
- **100BASE-TX-HD** PHY is running 100BASE-TX (100 mbps) in half-duplex mode.
- **100BASE-TX-FD** PHY is running 100BASE-TX (100 mbps) in full-duplex mode.
- **1000BASE-T-FD** PHY is running 1000BASE-T (1000 mbps) in full-duplex mode.
- **1000BASE-KX** PHY is 1000BASE-KX (1000 mbps) in full-duplex mode over backplane.
- **10GBASE-KX4** PHY is 10GBASE-KX4 (10 gbps) in full-duplex mode over backplane.
- **10GBASE-CR** PHY is 10GBASE-CR (10 gbps) in full-duplex mode.
- **10GBASE-SR** PHY is 10GBASE-SR (10 gbps) in full-duplex mode.
- **40GBASE-CR4** PHY is 40GBASE-CR4 (40 gbps) in full-duplex mode.
- **40GBASE-SR4** PHY is 40GBASE-SR4 (40 gbps) in full-duplex mode.
- **none** No PHY has been negotiated, no link.

Type

String.

The type of interface:

- **Ethernet** Indicates that it is on the base Ethernet MAC interface.

- **Tunnel** Indicates that it is on a encapsulation interface.
- **VLAN** Indicates that it is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.
- **Aggregate** Indicates that it is on an IEEE 802.3ad Link Aggregation composed of one of the base Ethernet MAC interface.
- **Other** Refers to network interfaces that you cannot configure.

MTU

Integer.

Identifies the maximum transmission unit (MTU) for the interface.

Aggregate

String.

Indicates the associated aggregate interface, if any.

LinkAddress

String.

Identifies the physical IP address for the interface. The default value for Ethernet interfaces is programmed in the Ethernet hardware.

Examples

The following example requests information about interfaces on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/LinkStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/LinkStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/LinkStatus"
    }
  },
  "LinkStatus": [
    {
      "Name": "lo",
      "Index": 1,
      "Status": "ok",
      "Mode": "none",
      "Type": "Other",
      "MTU": 65536,
      "Aggregate": "",
      "LinkAddress": "00:00:00:00:00:00"
    },
    {
      "Name": "mgt0",
      "Index": 4,
      "Status": "ok",
      "Mode": "1000BASE-T-FD",
      "Type": "Ethernet",
      "MTU": 1500,
      "Aggregate": "",
      "LinkAddress": "00:0b:ab:8a:56:ee"
    },
    {
      "Name": "mgt1",
      "Index": 7,
      "Status": "no-link",
      "Mode": "none",
      "Type": "Ethernet",
      "MTU": 1500,
      "Aggregate": "",
      "LinkAddress": "00:0b:ab:8a:56:ef"
    },
    {
      "Name": "sit0",
```

```

    "Index": 2,
    "Status": "no-link",
    "Mode": "none",
    "Type": "Other",
    "MTU": 1480,
    "Aggregate": "",
    "LinkAddress": "00:00:00:00:00:00"
  },
  {
    "Name": "eth10",
    "Index": 5,
    "Status": "no-link",
    "Mode": "none",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:84:23:59"
  },
  {
    "Name": "eth11",
    "Index": 8,
    "Status": "no-link",
    "Mode": "none",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:84:23:5a"
  },
  {
    "Name": "eth12",
    "Index": 9,
    "Status": "no-link",
    "Mode": "none",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:84:23:5b"
  },
  {
    "Name": "eth13",
    "Index": 11,
    "Status": "ok",
    "Mode": "1000BASE-T-FD",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:84:23:5c"
  },
  {
    "Name": "eth14",
    "Index": 12,
    "Status": "no-link",
    "Mode": "none",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:83:ec:cb"
  },
  {
    "Name": "eth15",
    "Index": 13,
    "Status": "no-link",
    "Mode": "none",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:83:ec:cc"
  },
  {
    "Name": "eth16",
    "Index": 14,
    "Status": "no-link",
    "Mode": "none",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:83:ec:cd"
  },
  {
    "Name": "eth17",
    "Index": 15,
    "Status": "no-link",

```

```

    "Mode": "none",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:83:ec:ce"
  },
  {
    "Name": "eth20",
    "Index": 6,
    "Status": "ok",
    "Mode": "10GBASE-CR",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:f0:7c:ba"
  },
  {
    "Name": "eth21",
    "Index": 10,
    "Status": "ok",
    "Mode": "10GBASE-SR",
    "Type": "Ethernet",
    "MTU": 1500,
    "Aggregate": "",
    "LinkAddress": "00:0b:ab:f0:7c:bb"
  },
  {
    "Name": "ip6tnl0",
    "Index": 3,
    "Status": "no-link",
    "Mode": "none",
    "Type": "Other",
    "MTU": 1452,
    "Aggregate": "",
    "LinkAddress": "00:00:00:00:00:00"
  }
]
}

```

[/mgmt/status/default/LoadBalancerStatus2](#)

You can use the HTTP GET method with the `/mgmt/status/default/LoadBalancerStatus2` resource to request status for any configured load balancer groups that the appliance is a member of.

GET

Use the HTTP GET method with the `/mgmt/status/default/LoadBalancerStatus2` resource to request status for any configured load balancer groups that the appliance is a member of.

- [Resource URL](#)
- [“Request headers” on page 1095](#)
- [Request body format](#)
- [“Security requirements” on page 1096](#)
- [Response status codes](#)
- [“Response headers” on page 1096](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/LoadBalancerStatus2
```

Where `port` is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about load balancer groups on the appliance. Each object contains the following fields:

Group

String.

The name of the load balancer group.

Server

String.

The IP address or host name of the load balancer member.

Port

Integer.

The member-specific target port or 0 if using the service-defined port.

Health

String.

The health status of a given load balancer member. Possible values are: `up`, `softdown`, or `down`.

Weight

Integer.

The weight given to this server. Weights are used in weighted load balancing decisions for algorithms such as `Weighted Round Robin` or `Weighted Least Outstanding Connections`.

AdminState

String.

The administrative state overrides the health of the server.

- **Enabled** This load balancer member is administratively enabled. The server participates in load balancing and health check decisions.
- **Disabled** This load balancer member is administratively disabled. The server is not included in any load balancing decisions.

Examples

The following example requests information about load balancer groups that the appliance belongs to.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/LoadBalancerStatus2
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/LoadBalancerStatus2"
    },
    "doc": {
      "href": "/mgmt/docs/status/LoadBalancerStatus2"
    }
  },
  "LoadBalancerStatus2": [
    {
      "Group": {
        "value": "pre_tend2",
        "href": "/mgmt/config/default/LoadBalancerGroup/pre_tend2"
      },
      "Server": "roundpen.com",
      "Port": 0,
      "Health": "up",
      "Weight": 1,
      "AdminState": "enabled"
    },
    {
      "Group": {
        "value": "pre_tend",
        "href": "/mgmt/config/default/LoadBalancerGroup/pre_tend"
      },
      "Server": "manege.com",
      "Port": 0,
      "Health": "up",
      "Weight": 1,
      "AdminState": "enabled"
    },
    {
      "Group": {
        "value": "pre_tend",
        "href": "/mgmt/config/default/LoadBalancerGroup/pre_tend"
      },
      "Server": "paddock.com",
      "Port": 0,
      "Health": "up",
      "Weight": 1,
      "AdminState": "enabled"
    }
  ]
}
```

[/mgmt/status/default/LogTargetStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/LogTargetStatus` resource to view the status for each log target on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/LogTargetStatus` resource to view the status for each log target on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1098](#)

- [Request body format](#)
- [“Security requirements” on page 1098](#)
- [Response status codes](#)
- [“Response headers” on page 1098](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/LogTargetStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about log targets on the appliance. Each object contains the following fields:

LogTarget

Value Name of the log target.

href URL of the log target.

Status

String.

The status of the log target. When the status indicates an error, check the error information column.

- **active** Normal operation state, no failures.
- **warning** Target is approaching failure conditions. Messages might be dropped.
- **failure** Target failed. No messages collected.
- **suspended** The target is busy because it is sending messages.

EventsProcessed

Integer.

The number of events that this log target processed.

EventsDropped

Integer.

The number of events that this log target dropped because there are too many pending events.

EventsPending

Integer.

The number of pending events for this log target. These events are waiting to be stored at the destination.

ErrorInfo

String.

Error information about the status for the log target.

RequestedMemory

String.

The requested memory for this log target. This measurement represents the high watermark of memory requested.

Examples

The following example requests information about log targets on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/LogTargetStatus
```

The following listing shows sections from the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/LogTargetStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/LogTargetStatus"
    }
  },
  "LogTargetStatus": [
    {
      "LogTarget": {
        "value": "default-log",
        "href": "/mgmt/config/default/LogTarget/default-log"
      },
      "Status": "active",
      "EventsProcessed": 895,
      "EventsDropped": 0,
      "EventsPending": 0,
      "ErrorInfo": "none",
      "RequestedMemory": 0
    },
    {
      "LogTarget": {
```

```

        "value": "logstash",
        "href": "/mgmt/config/default/LogTarget/logstash"
      },
      "Status": "active",
      "EventsProcessed": 1106,
      "EventsDropped": 0,
      "EventsPending": 0,
      "ErrorInfo": "",
      "RequestedMemory": 0
    ]
  }
}

```

[/mgmt/status/default/MQSystemResources](#)

You can use the HTTP GET method with the `/mgmt/status/default/MQSystemResources` resource to request status for information for IBM MQ resources on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/MQSystemResources` resource to request status for information for IBM MQ resources on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1100](#)
- [Request body format](#)
- [“Security requirements” on page 1100](#)
- [Response status codes](#)
- [“Response headers” on page 1101](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/MQSystemResources
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers**Content-Type**

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about IBM MQ resources on the appliance. Each object contains the following fields:

TotalStorage

Integer.

The total storage in megabytes available for IBM MQ data.

UsedStorage

Integer.

The amount of IBM MQ storage in use in megabytes.

TotalErrorsStorage

Integer.

The total storage in megabytes available for IBM MQ error logs.

UsedErrorsStorage

Integer.

The amount of IBM MQ error log storage in use in megabytes.

TotalTraceStorage

Integer.

The total storage in megabytes available for IBM MQ trace.

UsedTraceStorage

Integer.

The amount of IBM MQ trace storage in megabytes in use.

HAStatus

String.

Status of this IBM MQ appliance in the High Availability group, if any.

HAPartner

String.

Name and status of the partner appliance in the High Availability group, if any.

Examples

The following example requests information about the IBM MQ resources on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/MQSystemResources
```

The following listing shows an example received response:

```

{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/MQSystemResources"
    },
    "doc": {
      "href": "/mgmt/docs/status/MQSystemResources"
    }
  },
  "MQSystemResources": {
    "TotalStorage": 1142681,
    "UsedStorage": 328704,
    "TotalErrorsStorage": 16179,
    "UsedErrorsStorage": 176,
    "TotalTraceStorage": 32256,
    "UsedTraceStorage": 206,
    "HAStatus": "",
    "HAPartner": ""
  }
}

```

[/mgmt/status/default/MultipathStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/MultipathStatus` resource to request information about the routes available for volumes that are using multipath.

Note: This request is available only on M2000 and M2001 appliances.

GET

Use the HTTP GET method with the `/mgmt/status/default/MultipathStatus` resource to request information about the routes available for volumes that are using multipath.

Note: This request is available only on M2000 and M2001 appliances.

- [Resource URL](#)
- [“Request headers” on page 1102](#)
- [Request body format](#)
- [“Security requirements” on page 1103](#)
- [Response status codes](#)
- [“Response headers” on page 1103](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/MultipathStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about available multipath routes. Each object contains the following fields:

volumeName

String.

Name of the volume.

group

Integer.

Multipath group to which this path belongs.

groupPriority

Integer.

Numeric priority value for this group.

HBA

String.

The Host Bus Adapter associated with this path.

bus

Integer.

The number of the SCSI BUS associated with this path.

lun

Integer.

The SCSI Logical Unit number associated with this volume.

groupState

String.

Current state of the multipath path group. Possible states of a path group are:

- **active** If a path in the active path group is in the ready state, it is sent any I/O requests for the multipath device. At most one path group can be in the active state.
- **enabled** If no paths in the active path group are in the ready state, I/O requests are instead sent to the enabled path groups. Multiple path groups can be in the enabled state.
- **disabled** If no paths in the active or enabled path groups are in the ready state, I/O will be sent to the disabled path groups. Multiple path groups can be in the disabled state.

pathState

String.

Current state of the multipath path. Possible states of a path are:

- **active** The most recent I/O request for this path was successful.
- **failed** The most recent I/O request for this path failed.

deviceState

String.

Current state of the device associated with the path. Possible states of a device are:

- **ready** The path is up and is available for I/O requests.
- **faulty** The path is down and is not available for I/O requests.
- **ghost** The path is passive.
- **shaky** The path is up, but is temporarily not available for I/O requests.

Examples

The following example requests information about multipath routes.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/MultipathStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/MultipathStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/MultipathStatus"
    }
  },
  "MultipathStatus": [
    {
      "Volume": "FCVolume1",
      "GroupNumber": 1,
      "GroupPriority": 50,
      "HBA": "fch1",
      "Bus": 0,
      "LUN": 5,
      "GroupState": "active",
      "PathState": "active",
      "DeviceState": "ready"
    },
    {
      "Volume": "FCVolume1",
      "GroupNumber": 1,
      "GroupPriority": 50,
      "HBA": "fch2",
      "Bus": 0,
      "LUN": 5,
      "GroupState": "active",
      "PathState": "active",
      "DeviceState": "ready"
    },
    {
      "Volume": "FCVolume1",
      "GroupNumber": 2,
      "GroupPriority": 10,
```

```

    "HBA": "fch1",
    "Bus": 0,
    "LUN": 5,
    "GroupState": "enabled",
    "PathState": "active",
    "DeviceState": "ready"
  },
  {
    "Volume": "FCVolume1",
    "GroupNumber": 2,
    "GroupPriority": 10,
    "HBA": "fch2",
    "Bus": 0,
    "LUN": 5,
    "GroupState": "enabled",
    "PathState": "active",
    "DeviceState": "ready"
  }
]
}

```

[/mgmt/status/default/NDCacheStatus2](#)

You can use the HTTP GET method with the `/mgmt/status/default/NDCacheStatus2` resource to request status information about IPv6 Neighbor Discovery (ND) translations on all interfaces.

GET

Use the HTTP GET method with the `/mgmt/status/default/NDCacheStatus2` resource to request status information about IPv6 Neighbor Discovery (ND) translations on all interfaces. Only complete cache entries are returned.

- [Resource URL](#)
- [“Request headers” on page 1105](#)
- [Request body format](#)
- [“Security requirements” on page 1105](#)
- [Response status codes](#)
- [“Response headers” on page 1106](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/NDCacheStatus2
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about available ND translations. Each object contains the following fields:

IPType

String.

The version of the IP address: `ipv4` or `ipv6`:

- **unknown** An unknown address type.
- **ipv4** An IPv4 address.
- **ipv6** An IPv6 address.
- **ipv4z** A nonglobal IPv4 address with a zone index.
- **ipv6z** A nonglobal IPv6 address with a zone index.
- **dns** A DNS domain name

IPAddress

String.

The IPv6 address that corresponds to the MAC address.

PrefixLength

Integer.

The prefix length of the IP address that represents the configured netmask.

HardwareAddress

String.

The Host Bus Adapter associated with this path.

Interface

String.

The name of the interface.

InterfaceIndex

Integer.

The index of the interface.

InterfaceType

String.

The interface type:

- **Ethernet** Indicates that it is on the base Ethernet MAC interface.
- **Tunnel** Indicates that it is on a encapsulation interface.
- **VLAN** Indicates that it is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.
- **Aggregate** Indicates that it is on an IEEE 802.3ad Link Aggregation composed of one of the base Ethernet MAC interface.
- **Other** Refers to network interfaces that you cannot configure.

State

String.

The lifecycle state for the entry:

- **incomplete** The cache entry is in the incomplete state. It is inaccessible.
- **reachable** The cache entry is in the reachable state. It is accessible.
- **stale** The cache entry is in the stale state. It is potentially accessible.
- **delay** The cache entry is in the delay state. It is inaccessible.
- **probe** The cache entry is in the probe state. It is inaccessible.
- **failed** The cache entry is in the failed state. It is inaccessible.
- **noarp** The cache entry is in the noarp state. It is inaccessible.
- **permanent** The cache entry is in the permanent state. It is always accessible.

Examples

The following example requests information about ND translations.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/NDCacheStatus2
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/NDCacheStatus2"
    },
    "doc": {
      "href": "/mgmt/docs/status/NDCacheStatus2"
    }
  },
  "result": "No status retrieved."
}
```

[/mgmt/status/default/NetworkInterfaceStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/NetworkInterfaceStatus` resource to request status information about the network interfaces on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/NetworkInterfaceStatus` resource to request information about the network interfaces on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1108](#)
- [Request body format](#)
- [“Security requirements” on page 1108](#)
- [Response status codes](#)
- [“Response headers” on page 1108](#)

- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/NetworkInterfaceStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated”](#) on page 423 and [“User authorization, credential mapping, and access profiles”](#) on page 423.

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about available network interfaces. Each object contains the following fields:

InterfaceIndex

Index.

The value of the SNMP `ifIndex` column in the `ifTable` for the interface.

InterfaceType

String.

- **Ethernet** Indicates that it is on the base Ethernet MAC interface.

- **Tunnel** Indicates that it is on an encapsulation interface.
- **VLAN** Indicates that it is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.
- **Aggregate** Indicates that it is on an IEEE 802.3ad Link Aggregation composed of one of the base Ethernet MAC interface.
- **Other** Refers to network interfaces that you cannot configure.

Name

String.

The configuration name of the interface.

AdminStatus

String.

The configured administrative state of the interface.

- **Enabled** The administrative state of the interface is enabled.
- **Disabled** The administrative state of the interface is disabled.

OperStatus

The current operational state of the interface.

- **up** The interface is operationally up.
- **down** The interface is operationally down.
- **testing** Some testing mode. After the test completes, the status changes as appropriate.
- **unknown** Cannot determine the operational state of the interface.
- **dormant** The interface is waiting for a connection.
- **notPresent** The interface is not present.
- **LowerLayerDown** The required lower layer is down.

IPType

String.

The IP address version: ipv4 or ipv6.

- **unknown** An unknown address type.
- **ipv4** An IPv4 address.
- **ipv6** An IPv6 address.
- **ipv4z** A nonglobal IPv4 address with a zone index.
- **ipv6z** A nonglobal IPv6 address with a zone index.
- **dns** A DNS domain name

IP

String.

The primary IP address for the interface.

PrefixLength

String.

The prefix length of the IP address that represents the configured netmask.

MACAddress

String.

The current MAC address of the interface. The address can be the pre-programmed one, overridden by your configuration, or overridden by standby control if the interface is the active member of the standby group.

MTU

Integer.

Maximum transmission unit for the interface.

RxHCBytes

Integer.

The amount of data successfully received on the interface, which includes MAC framing overhead.

RxHCPackets

Integer.

The number of packets successfully received on the interface that were passed up to the network layer for processing.

RxErrors2

Integer.

The number of packets that could not be received due to errors in the packet or in the hardware.

RxDrops2

Integer.

The number of received packets that were not in error, but were not passed up to the network layer due to resource constraints.

TxHCBytes

Integer.

The amount of data successfully transmitted on the interface, which includes MAC framing overhead.

TxHCPackets

Integer.

The number of packets successfully transmitted on the interface.

TxErrors2

Integer.

The number of packets that were not successfully transmitted due to errors on the network or in the hardware.

TxDrops2

Integer.

The number of packets that were not transmitted because the network layer was generating packets faster than the physical network could accept them.

Examples

The following example requests information about network interfaces.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/NetworkInterfaceStatus
```

The following listing shows sections from an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/NetworkInterfaceStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/NetworkInterfaceStatus"
    }
  },
  "NetworkInterfaceStatus": [
    {
      "InterfaceIndex": 1,
      "InterfaceType": "Other",
      "Name": "lo",
      "AdminStatus": "up",
      "OperStatus": "up",
      "IPType": "ipv4",
      "IP": "127.0.0.1",
      "PrefixLength": 8,
      "MACAddress": "00:00:00:00:00:00",
      "MTU": 65536,
    }
  ]
}
```

```

    "RxHCBytes": 1869015199,
    "RxHCPackets": 14005952,
    "RxErrors2": 0,
    "RxDrops2": 0,
    "TxHCBytes": 1869015199,
    "TxHCPackets": 14005952,
    "TxErrors2": 0,
    "TxDrops2": 0
  },
  {
    "InterfaceIndex": 2,
    "InterfaceType": "Other",
    "Name": "sit0",
    "AdminStatus": "down",
    "OperStatus": "lowerLayerDown",
    "IPType": "ipv4",
    "IP": "",
    "PrefixLength": 0,
    "MACAddress": "00:00:00:00:00:00",
    "MTU": 1480,
    "RxHCBytes": 0,
    "RxHCPackets": 0,
    "RxErrors2": 0,
    "RxDrops2": 0,
    "TxHCBytes": 0,
    "TxHCPackets": 0,
    "TxErrors2": 0,
    "TxDrops2": 0
  },
  ...
  {
    "InterfaceIndex": 15,
    "InterfaceType": "Ethernet",
    "Name": "eth17",
    "AdminStatus": "down",
    "OperStatus": "lowerLayerDown",
    "IPType": "ipv4",
    "IP": "",
    "PrefixLength": 0,
    "MACAddress": "00:0b:ab:83:ec:ce",
    "MTU": 1500,
    "RxHCBytes": 0,
    "RxHCPackets": 0,
    "RxErrors2": 0,
    "RxDrops2": 0,
    "TxHCBytes": 0,
    "TxHCPackets": 0,
    "TxErrors2": 0,
    "TxDrops2": 0
  }
]
}

```

[/mgmt/status/default/NFSMountStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/NFSMountStatus` resource to request status information about NFS mounts on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/NFSMountStatus` resource to request status information about NFS mounts on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1112](#)
- [Request body format](#)
- [“Security requirements” on page 1112](#)
- [Response status codes](#)
- [“Response headers” on page 1112](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/NFSMountStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about NFS mounts. Each object contains the following fields:

Type

String.

The type of mount service this NFS mount provides.

- Static URL
- Dynamic URL

Name

String.

The configuration identifier for this NFS mount.

Remote

String.

The remote NFS mount point for this mount.

State

String.

The state of this NFS mount. If the state is up, files can be opened over this mount.

- **new** The mount is newly created. Mounting has not yet started.
- **mounting** The mount process is in progress.
- **up** The mount process is complete and healthy. Files can be opened.
- **down** The mount process is complete and healthy. Files can be opened.
- **unmounting** The unmount process is in progress.
- **unmountingForRemount** The unmount process is in progress to change mount parameters. When the unmount process is complete, will remount.
- **unmounted** The unmount process is complete.
- **defective** The mount process cannot complete because of a corrupt state or because of the configuration.

Version

String.

The version of the NFS protocol being used for this mount.

Address

String.

The IP address of the interface of the NFS server that NFS RPC traffic is being sent to, or was last sent to.

Transport

String.

The transport protocol being used under the RPC protocol for this NFS mount.

- **TCP** Uses TCP. TCP is more reliable and allows larger read and write operations than UDP. If selected and not available, falls back to UDP.
- **UDP** Uses UDP. Some early NFS Version 2 servers support only UDP.

Port

Integer.

The UDP or TCP port number that RPC data for this NFS mount is sent to.

FilesRead

Integer.

The number of files that were opened for reading over this NFS mount.

FilesWritten

Integer.

The number of files that were opened for write (create) over this NFS mount.

FilesOpen

Integer.

The number of files that are currently open for read or write (create) over this NFS mount.

Examples

The following example requests information about NFS mounts.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/NFSMountStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/NFSMountStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/NFSMountStatus"
    }
  },
  "result": "No status retrieved."
}
```

[/mgmt/status/default/NTPRefreshStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/NTPRefreshStatus` resource to request status information about the most recent NTP refresh.

GET

Use the HTTP GET method with the `/mgmt/status/default/NTPRefreshStatus` resource to request status information about the most recent NTP refresh.

- [Resource URL](#)
- [“Request headers” on page 1114](#)
- [Request body format](#)
- [“Security requirements” on page 1114](#)
- [Response status codes](#)
- [“Response headers” on page 1115](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/NTPRefreshStatus
```

Where `port` is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers**Content-Type**

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the most recent NTP refresh. Each object contains the following fields:

LastRefreshIndex

String.

IP address of the last server contacted as an NTP server.

LastRefreshResult

String.

Result of the last server contact.

LastRefreshTime

String.

Time after the last refresh.

Time

String.

The current local time. This time is obtained from the timezone settings established on the Time Settings object and with the date and time settings set on the System Control page.

Examples

The following example requests information about the most recent NTP refresh.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/NTPRefreshStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/NTPRefreshStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/NTPRefreshStatus"
    }
  },
  "NTPRefreshStatus": {
    "LastRefreshIndex": "9.180.171.67",
    "LastRefreshResult": "Success!",
    "LastRefreshTime": "Fri Feb 8 14:41:56 2019",
    "Time": "Fri Feb 8 14:49:22 2019"
  }
}
```

```
}  
}
```

[/mgmt/status/default/ObjectStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/ObjectStatus` resource to request information about the status of configuration objects on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/ObjectStatus` resource to request status information about the configuration objects on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1116](#)
- [Request body format](#)
- [“Security requirements” on page 1116](#)
- [Response status codes](#)
- [“Response headers” on page 1117](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/ObjectStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about configuration objects on the appliance. Each object contains the following fields:

Class

String.

The type of configuration object.

OpState

String.

The operational state of the object. Can be up or down.

AdminState

String.

The administrative state of the configuration object. Can be enabled or disabled.

Name

String.

The name of the configuration object.

EventCode

String.

The event code associated with the object.

ErrorCode

String.

Any text that is associated with the EventCode.

ConfigState

String.

The configuration status. Can be one of the following:

- **Invalid** Invalid configuration.
- **Saved** Persisted configuration.
- **New** New configuration.
- **Modified** Modified configuration.
- **Deleted** Deleted configuration.
- **External** External configuration.

Examples

The following example requests information about configuration objects.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/ObjectStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/ObjectStatus"
    }
  }
}
```

```

    },
    "doc": {
      "href": "/mgmt/docs/status/ObjectStatus"
    }
  },
  "ObjectStatus": [
    {
      "Class": "DNSNameService",
      "OpState": "up",
      "AdminState": "enabled",
      "Name": "dns",
      "EventCode": "0x00000000",
      "ErrorCode": "",
      "ConfigState": "saved"
    },
    {
      "Class": "EthernetInterface",
      "OpState": "down",
      "AdminState": "disabled",
      "Name": "eth0",
      "EventCode": "0x0034000d",
      "ErrorCode": "Object is disabled",
      "ConfigState": "saved"
    },
    {
      "Class": "EthernetInterface",
      "OpState": "down",
      "AdminState": "disabled",
      "Name": "eth1",
      "EventCode": "0x0034000d",
      "ErrorCode": "Object is disabled",
      "ConfigState": "saved"
    },
    ...
    {
      "Class": "AuditLog",
      "OpState": "up",
      "AdminState": "enabled",
      "Name": "audit-log-settings",
      "EventCode": "0x00000000",
      "ErrorCode": "",
      "ConfigState": "saved"
    },
    {
      "Class": "RestMgmtInterface",
      "OpState": "up",
      "AdminState": "enabled",
      "Name": "rest-mgmt",
      "EventCode": "0x00000000",
      "ErrorCode": "",
      "ConfigState": "saved"
    },
    {
      "Class": "SSHServerProfile",
      "OpState": "up",
      "AdminState": "enabled",
      "Name": "sshserverprofile",
      "EventCode": "0x00000000",
      "ErrorCode": "",
      "ConfigState": "saved"
    }
  ]
}

```

[/mgmt/status/default/OtherSensors](#)

You can use the HTTP GET method with the `/mgmt/status/default/OtherSensors` resource to request status information about the sensors on the appliance that have true or false values.

GET

Use the HTTP GET method with the `/mgmt/status/default/OtherSensors` resource to request status information about the sensors on the appliance that have true or false values.

- [Resource URL](#)
- [“Request headers” on page 1119](#)
- [Request body format](#)

- [“Security requirements” on page 1119](#)
- [Response status codes](#)
- [“Response headers” on page 1119](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/OtherSensors
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about appliance sensors. Each object contains the following fields:

Name

String.

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.

Value

String.

The value of the sensor. The value is either true or false.

ReadingStatus

String.

The current 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.
- **noReading** No reading of the sensor is available at this time. An internal hardware or software error might occur. Contact IBM Support.
- **invalid** The appliance software requested a sensor reading by using an invalid sensor identifier. An internal software error might occur. Contact IBM Support.

Examples

The following example requests information about appliance sensors.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/OtherSensors
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/OtherSensors"
    },
    "doc": {
      "href": "/mgmt/docs/status/OtherSensors"
    }
  },
  "OtherSensors": [
    {
      "Name": "Intrusion Detected",
      "Value": "false",
      "ReadingStatus": "ok"
    },
    {
      "Name": "Power Supply 1 AC Lost",
      "Value": "false",
      "ReadingStatus": "ok"
    },
    {
      "Name": "Power Supply 1 Output Failure",
      "Value": "true",
      "ReadingStatus": "failure"
    },
    {
      "Name": "Power Supply 2 AC Lost",
      "Value": "false",
      "ReadingStatus": "ok"
    },
    {
      "Name": "Power Supply 2 Output Failure",
      "Value": "false",
      "ReadingStatus": "ok"
    }
  ]
}
```

[/mgmt/status/default/PCIBus](#)

You can use the HTTP GET method with the `/mgmt/status/default/PCIBus` resource to retrieve a list of the PCI functions on the system. The functions are identified by domain, bus, device, and function.

GET

Use the HTTP GET method with the `/mgmt/status/default/PCIBus` resource to request status information about PCI functions on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1121](#)
- [Request body format](#)
- [“Security requirements” on page 1121](#)
- [Response status codes](#)
- [“Response headers” on page 1122](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/PCIBus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of application/json.

Response body format

The response is in JSON format. The response includes information about PCI functions. Each object contains the following fields:

Address

String.

The PCI address of this function. The address is a series of hexadecimal numbers, in the format xxxx:xx:xx.x, with each group being domain, bus, target, and function in turn. So 0000:01:02.3 is domain 0, bus 1, target 2, and function 3.

VendorID

String.

The vendor identification of this PCI function.

DeviceID

String.

The device identification of this PCI function. Each vendor assigns their own device identifications.

Class

String.

The class of this PCI function. Class values are assigned by the PCI SIG.

Revision

String.

The Revision of this PCI function.

InterruptLine

String.

The interrupt line assigned to this PCI function, if any. If none is assigned, no value is displayed.

Examples

The following example requests information about appliance PCI functions.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/PCIBus
```

The following listing shows an extract from an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/PCIBus"
    },
    "doc": {
      "href": "/mgmt/docs/status/PCIBus"
    }
  },
  "PCIBus": [
    {
      "Address": "0000:00:00.0",
      "VendorID": "0x8086",
      "DeviceID": "0x0E00",
      "Class": "0x0600",
      "Revision": 4,
      "InterruptLine": ""
    },
    {
      "Address": "0000:00:01.0",
      "VendorID": "0x8086",
      "DeviceID": "0x0E02",

```

```

    "Class": "0x0604",
    "Revision": 4,
    "InterruptLine": 11
  },
  {
    "Address": "0000:00:02.0",
    "VendorID": "0x8086",
    "DeviceID": "0x0E04",
    "Class": "0x0604",
    "Revision": 4,
    "InterruptLine": 11
  },
  ...
  {
    "Address": "0000:ff:16.1",
    "VendorID": "0x8086",
    "DeviceID": "0x0EC9",
    "Class": "0x0880",
    "Revision": 4,
    "InterruptLine": ""
  },
  {
    "Address": "0000:ff:16.2",
    "VendorID": "0x8086",
    "DeviceID": "0x0ECA",
    "Class": "0x0880",
    "Revision": 4,
    "InterruptLine": ""
  }
]
}

```

[/mgmt/status/default/QueueManagersStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/QueueManagersStatus` resource to retrieve detailed status of available queue managers.

GET

Use the HTTP GET method with the `/mgmt/status/default/QueueManagersStatus` resource to request detailed status of available queue managers.

- [Resource URL](#)
- [“Request headers” on page 1123](#)
- [Request body format](#)
- [“Security requirements” on page 1124](#)
- [Response status codes](#)
- [“Response headers” on page 1124](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/QueueManagersStatus
```

Where `port` is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about queue manager status. Each object contains the following fields:

Name

String.

Queue manager name.

Status

String.

Run time status of the queue manager.

CpuUsage

String.

The instantaneous CPU usage by the queue manager as a percentage of the CPU load.

UsedMemory

Integer.

The amount of memory that is currently in use by the queue manager.

UsedFs

Integer.

The amount of file system that is currently in use by the queue manager.

TotalFs

Integer.

The amount of file system allocated for the queue manager.

HaRole

String.

HA role of the queue manager.

HaStatus

String.

HA status of the queue manager.

DrRole

String.

DR role of the queue manager.

DrStatus

String.

DR status of the queue manager.

Examples

The following example requests information about queue managers on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/QueueManagersStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/QueueManagersStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/QueueManagersStatus"
    }
  },
  "QueueManagersStatus": [
    {
      "Name": "QM1",
      "Status": "Running",
      "CpuUsage": "0.00",
      "UsedMemory": 187,
      "UsedFs": 245,
      "TotalFs": 64512,
      "HaRole": "NA",
      "HaStatus": "NA",
      "DrRole": "NA",
      "DrStatus": "NA"
    },
    {
      "Name": "QM2",
      "Status": "Ended immediately",
      "CpuUsage": "0.00",
      "UsedMemory": 0,
      "UsedFs": 245,
      "TotalFs": 64512,
      "HaRole": "NA",
      "HaStatus": "NA",
      "DrRole": "NA",
      "DrStatus": "NA"
    },
    {
      "Name": "QM3",
      "Status": "Running",
      "CpuUsage": "0.00",
      "UsedMemory": 187,
      "UsedFs": 245,
      "TotalFs": 64512,
      "HaRole": "NA",
      "HaStatus": "NA",
      "DrRole": "NA",
      "DrStatus": "NA"
    }
  ]
}
```

[/mgmt/status/default/RaidArrayStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/RaidArrayStatus` resource to retrieve details about the status of the RAID array.

GET

Use the HTTP GET method with the `/mgmt/status/default/RaidArrayStatus` resource to retrieve details about the status of the RAID array.

- [Resource URL](#)
- [“Request headers” on page 1126](#)
- [Request body format](#)
- [“Security requirements” on page 1126](#)
- [Response status codes](#)
- [“Response headers” on page 1127](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/RaidArrayStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the raid array on the appliance and contains the following fields:

ControllerID

Integer.

The reference number of the RAID card. The value is always 1.

ArrayID

Integer.

The reference number of the this array. Numbering starts with 1.

LogicalDriveID

Integer.

The identifier of the logical driver of which this array is part.

RaidLevel

String.

The RAID level of this array configuration:

- **raidLevel0** The configuration is RAID 0. Data is written without redundancy or fault tolerance.
- **raidLevel1** The configuration is RAID 1. Data is written redundantly.
- **raidLevel1E** The configuration is RAID 1E. Data is written redundantly.
- **raidLevel5** The configuration is RAID 5. Data is accessible after a failure to a single drive.
- **raidLevel6** The configuration is RAID 6. Data is accessible after a failure to two drives.
- **raidLevel10** The configuration is RAID 10, which is RAID 1 with spanning.
- **raidLevel50** The configuration is RAID 50, which is RAID 5 with spanning.
- **raidLevel60** The configuration is RAID 60, which is RAID 6 with spanning.
- **undefined** Cannot detect the RAID level. Check your configuration.

NumPhysicalDrives

Integer.

The number of the physical drives for this array.

TotalSize

Integer.

The normalized size of this array in megabytes. The value is rounded down to an even multiple, which allows you to swap drives of the same nominal size but might not be the same raw size.

Examples

The following example requests information about the raid array on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/RaidArrayStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/RaidArrayStatus"
    },
    "doc": {
```

```

    "href": "/mgmt/docs/status/RaidArrayStatus"
  }
},
"RaidArrayStatus": {
  "ControllerID": 1,
  "ArrayID": 1,
  "LogicalDriveID": 1,
  "RaidLevel": "raidLevel1",
  "NumPhysicalDrives": 2,
  "TotalSize": 1143455
}
}
}

```

[/mgmt/status/default/RaidBatteryModuleStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/RaidBatteryModuleStatus` resource to retrieve details about the battery backup unit (BBU) of the RAID controller. The BBU protects against the loss of cached data in the event of a power failure.

GET

Use the HTTP GET method with the `/mgmt/status/default/RaidBatteryModuleStatus` resource to retrieve details about the battery backup unit (BBU) of the RAID controller.

- [Resource URL](#)
- [“Request headers” on page 1128](#)
- [Request body format](#)
- [“Security requirements” on page 1128](#)
- [Response status codes](#)
- [“Response headers” on page 1129](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/RaidBatteryModuleStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers**Content-Type**

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the RAID battery backup unit on the appliance. The response contains the following fields:

ControllerID

Integer.

The reference number of the RAID controller. The value is always 1.

BatteryType

String.

The type of the BBU.

Serial

Integer.

The serial number of the BBU.

Name

String.

The product name of the BBU.

Status

String.

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 of the battery failed because of errors in the learning cycle.
- **learnCycleRequested** A request is submitted for a learning cycle.
- **learnCycleTimeout** The battery learning cycle has timed out.
- **packMissing** The battery is unplugged.
- **temperatureHigh** The battery temperature is high.
- **voltageLow** The battery voltage is low.
- **periodicLearnRequired** A request is submitted for a learning cycle to perform battery maintenance.
- **remainingCapacityLow** The remaining capacity of the battery is low.
- **replacePack** The battery needs to be replaced.
- **normal** The battery is in a normal state.
- **undefined** The battery state is undefined.

- **transparentLearn** The transparent learn mode of the battery is enabled.

Voltage

Integer.

The actual voltage of the battery in millivolts.

Current

Integer.

The current that flows through the battery terminals in milliamperes.

Temperature

Integer.

The temperature of the battery in degrees Celsius.

DesignCapacity

Integer.

The designed capacity of the battery in milliampere-hour.

DesignVoltage

Integer.

The designed voltage of the battery in millivolts.

Examples

The following example requests information about the RAID battery backup unit on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/RaidBatteryModuleStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/RaidBatteryModuleStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/RaidBatteryModuleStatus"
    }
  },
  "RaidBatteryModuleStatus": {
    "ControllerID": 1,
    "BatteryType": "CVPM02",
    "Serial": 3707,
    "Name": "CVPM02",
    "Status": "normal",
    "Voltage": 9374,
    "Current": 0,
    "Temperature": 23,
    "DesignCapacity": 306,
    "DesignVoltage": 9500
  }
}
```

[/mgmt/status/default/RaidLogicalDriveStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/RaidLogicalDriveStatus` resource to retrieve details about the RAID logical drive.

GET

Use the HTTP GET method with the `/mgmt/status/default/RaidLogicalDriveStatus` resource to retrieve details about the RAID logical drive.

- [Resource URL](#)
- [“Request headers” on page 1131](#)
- [Request body format](#)

- [“Security requirements” on page 1131](#)
- [Response status codes](#)
- [“Response headers” on page 1131](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/RaidLogicalDriveStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the RAID logical drive. The response contains the following fields:

ControllerID

Integer.

The reference number of the RAID card. The value is always 1.

LogicalDriveID

Integer.

The reference number of this logical drive.

LogicalDriveName

String.

The name of the logical drive. The value is always 'raid0'.

RaidLevel

String.

The configuration of the RAID level for the logical disk in the RAID volume.

- **raidLevel0** The configuration is RAID 0. Data is written without redundancy or fault tolerance.
- **raidLevel1** The configuration is RAID 1. Data is written redundantly.
- **raidLevel1E** The configuration is RAID 1E. Data is written redundantly.
- **raidLevel5** The configuration is RAID 5. Data is accessible after a failure to a single drive.
- **raidLevel6** The configuration is RAID 6. Data is accessible after a failure to two drives.
- **raidLevel10** The configuration is RAID 10, which is RAID 1 with spanning.
- **raidLevel50** The configuration is RAID 50, which is RAID 5 with spanning.
- **raidLevel60** The configuration is RAID 60, which is RAID 6 with spanning.
- **undefined** Cannot detect the RAID level. Check your configuration.

NumPhysicalDrives

Integer.

The number of physical drives in the logical drive.

State

String.

The state of a logical disk in a RAID volume.

- **offline** The logical drive is offline. Check the physical drive status.
- **partiallyDegraded** One of physical drives in the logical drive is damaged. Replace the physical drive.
- **degraded** One of physical drive in the logical drive is damaged. Replace the physical drive.
- **optimal** The logical drive is functional.

InitProgress

String.

The progress of the initialization of the logical drive.

- **noInitInProgress** There is no initialization in progress for this logical drive.
- **quickInitInProgress** The logical drive is in the quick initialization state.
- **fullInitInProgress** The logical drive is in the full initialization state.
- **unknown** Cannot determine the initialization progress for this logical drive.

CurrReadPolicy

String.

- **noReadAhead** The current read policy is No Read Ahead.
- **adaptiveReadAhead** The current read policy is Adaptive Read Ahead.

CurrWritePolicy

String.

- **writeThrough** The current write policy is Write Through.
- **writeBack** The current write policy is Write Back.

- **writeCachingOkIfBadBBU** The current write policy is Write caching ok if bad BBU.

CurrCachePolicy

String.

The cache policy in effect.

- **cachedIO** The current cache policy is Cached IO.
- **directIO** The current cache policy is Direct IO.

AccessPolicy

String.

The access policy in effect for the logical drive.

- **readWrite** The access policy is Read Write.
- **readOnly** The access policy is Read Only.
- **blocked** The access policy is Blocked.
- **unknown** Cannot determine the access policy of a logical drive.

BadBlockExists

String.

An indicator of whether there are bad blocks in the logical drive.

- **no** There is no bad block in the logical drive.
- **yes** There is at least one bad block in the logical drive.

TotalSize

Integer.

The total size of the logical drive in megabytes.

Examples

The following example requests information about the RAID logical drive on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/RaidLogicalDriveStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/RaidLogicalDriveStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/RaidLogicalDriveStatus"
    }
  },
  "RaidLogicalDriveStatus": {
    "ControllerID": 1,
    "LogicalDriveID": 1,
    "LogicalDriveName": "raid0",
    "RaidLevel": "raidLevel1",
    "NumPhysicalDrives": 2,
    "State": "optimal",
    "InitProgress": "noInitInProgress",
    "CurrReadPolicy": "adaptiveReadAhead",
    "CurrWritePolicy": "writeBack",
    "CurrCachePolicy": "directIO",
    "AccessPolicy": "readWrite",
    "BadBlockExists": "noBadBlock",
    "TotalSize": 1143455
  }
}
```

[/mgmt/status/default/RaidPhysicalDriveStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/RaidPhysicalDriveStatus` resource to retrieve details about the RAID physical drives.

GET

Use the HTTP GET method with the `/mgmt/status/default/RaidPhysicalDriveStatus` resource to retrieve details about the RAID physical drives.

- [Resource URL](#)
- [“Request headers” on page 1134](#)
- [Request body format](#)
- [“Security requirements” on page 1134](#)
- [Response status codes](#)
- [“Response headers” on page 1135](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/RaidPhysicalDriveStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the RAID physical drives. Each object contains the following fields:

ControllerID

Integer.

The reference number of the RAID card. The value is always 1.

DeviceID

Integer.

The device identifier of the physical drive. Numbering starts with 1.

ArrayID

Integer.

The RAID array identifier to which this physical drive joins.

LogicalDriveID

Integer.

The logical drive identifier of which this physical drive is a part.

LogicalDriveName

String.

The name of the logical drive. The value is always 'raid0'.

Position

String.

The location of the physical drive from the front panel point-of-view of the appliance.

- **HDD0** The HDD0 drive is in the slot marked as 0 on the front panel.
- **HDD1** The HDD1 drive is in the slot marked as 1 on the front panel.
- **HDD2** The HDD2 drive is in the slot marked as 2 on the front panel.
- **HDD3** The HDD3 drive is in the slot marked as 3 on the front panel.
- **B5S1** The drive is in the first slot of the bay marked 5.
- **B5S2** The drive is in the second slot of the bay marked 5.
- **B7S1** The drive is in the first slot of the bay marked 7.
- **B7S2** The drive is in the second slot of the bay marked 7.
- **undefined** Cannot detect the location of the drive.

State

String.

The current state of the physical drive.

- **unconfiguredGood** The physical drive is good and available to be added in a RAID volume.
- **unconfiguredGoodForeign** The physical drive is good, but the configuration was created by another system.
- **unconfiguredBad** The physical drive is bad and cannot be added in a RAID volume. Replace this drive.
- **unconfiguredBadForeign** The physical drive is bad and the configuration was created by another system. Replaced this drive.
- **hotSpare** The physical drive is a hot spare disk.
- **offline** The physical drive is configured and good, but the data is invalid.

- **failed** The physical drive is configured but bad, all data is invalid and should be replaced.
- **rebuild** The physical drive is configured and rebuilding.
- **online** The physical drive is configured and good.
- **copyback** The physical drive is getting copied.
- **system** The physical drive is exposed and controlled by host.
- **undefined** The physical drive state is undefined.

ProgressPercent

Integer.

The current progress percentage of the operation on the physical drive. Operations can be rebuild, copyback, patrol, or clear.

RawSize

Integer.

The exact size of the drive in megabytes.

CoercedSize

Integer.

The normalized size in megabytes. The value is rounded down to an even multiple, which allows you to swap drives of the same nominal size but might not be the same raw size.

InterfaceType

String.

The interface type of the drive.

- **unknown** The drive type is unknown.
- **parallelSCSI** The drive type is parallel Small Computer System Interface (SCSI).
- **sas** The drive type is Serial Attached SCSI (SAS).
- **sata** The drive type is Serial Advanced Technology Attachment (SATA).
- **fiberChannel** The drive type is Fiber Channel.
- **undefined** The drive type is not defined.

InterfaceSpeed

String.

The SAS interface speed negotiated between the RAID controller and the physical drive.

- **serial1point5GigabitsPerSecond** The speed is 1.5 Gbps.
- **serial3point0GigabitsPerSecond** The speed is 3.0 Gbps.
- **serial6point0GigabitsPerSecond** The speed is 6.0 Gbps.
- **notApplicable** Cannot detect the interface speed.

SASaddress

String.

The SAS address of a physical drive.

VendorID

String.

The vendor identification string for the hard disk drive.

ProductID

String.

The product identification string for the hard disk drive.

Revision

String.

The revision string of the hard disk drive from the vendor.

SpecificInfo

String.

The vendor specific information string that is read from the hard drive. Generally, the identifier is unique for each physical drive.

Failure

String.

If the hard disk failure state shows Yes, replace this drive as soon as possible to avoid possible data loss.

- **yes** The physical drive has an error.
- **no** The physical drive has no errors.

Temperature

String.

The temperature of the hard disk drive.

Examples

The following example requests information about the RAID physical drives on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/RaidPhysicalDriveStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/RaidPhysicalDriveStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/RaidPhysicalDriveStatus"
    }
  },
  "RaidPhysicalDriveStatus": [
    {
      "ControllerID": 1,
      "DeviceID": 8,
      "ArrayID": 1,
      "LogicalDriveID": 1,
      "LogicalDriveName": "raid0",
      "Position": "HDD2",
      "State": "online",
      "ProgressPercent": 0,
      "RawSize": 1144641,
      "CoercedSize": 1143455,
      "InterfaceType": "sas",
      "InterfaceSpeed": "serial6point0GigabitsPerSecond",
      "SASaddress": "5000cca01daa4e050000000000000000",
      "VendorID": "IBM-ESXS",
      "ProductID": "HUC101212CSS60",
      "Revision": "J2M8",
      "SpecificInfo": "KZK0NHUJ2M8J2M8J2M8",
      "Failure": "No",
      "Temperature": "33 / 91"
    },
    {
      "ControllerID": 1,
      "DeviceID": 9,
      "ArrayID": 1,
      "LogicalDriveID": 1,
      "LogicalDriveName": "raid0",
      "Position": "HDD1",
      "State": "online",
      "ProgressPercent": 0,
      "RawSize": 1144641,
      "CoercedSize": 1143455,
      "InterfaceType": "sas",
      "InterfaceSpeed": "serial6point0GigabitsPerSecond",
      "SASaddress": "5000cca01daa39990000000000000000",
      "VendorID": "IBM-ESXS",
```

```
"ProductID": "HUC101212CSS60",
"Revision": "J2M8",
"SpecificInfo": "KZK0M4NHJ2M8J2M8J2M8",
"Failure": "No",
"Temperature": "33 / 91"
  }
]
```

[/mgmt/status/default/RaidSsdStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/RaidSsdStatus` resource to retrieve details about the status of the SSD disks on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/RaidSsdStatus` resource to retrieve details about the SSD disks on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1138](#)
- [Request body format](#)
- [“Security requirements” on page 1138](#)
- [Response status codes](#)
- [“Response headers” on page 1139](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/RaidSsdStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the SSD drives. Each object contains the following fields:

DiskNumber

Integer.

The number of the disk in the RAID volume. Numbering starts with 1.

SN

String.

The serial number string for the drive.

TotalWritten

Integer.

The total data in gigabytes written to the drive since manufacture.

lifeLeft

String.

Estimate of the remaining drive lifetime.

Examples

The following example requests information about the SSD disks on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/RaidSsdStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/RaidSsdStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/RaidSsdStatus"
    }
  },
  "RaidSsdStatus": [
    {
      "DiskNumber": 1,
      "SN": "ZAM15AR3",
      "TotalWritten": 3083,
      "LifeLeft": 100
    },
    {
      "DiskNumber": 2,
      "SN": "ZAM15ASJ",
      "TotalWritten": 3083,
      "LifeLeft": 100
    },
    {
      "DiskNumber": 3,
      "SN": "ZAM15AR6",
      "TotalWritten": 3080,
      "LifeLeft": 100
    }
  ]
}
```

```
    },
    {
      "DiskNumber": 4,
      "SN": "ZAM15AN5",
      "TotalWritten": 3080,
      "LifeLeft": 100
    }
  ]
}
```

[/mgmt/status/default/RoutingStatus3](#)

You can use the HTTP GET method with the `/mgmt/status/default/RoutingStatus3` resource to retrieve routing table details. This table describes the IP routes on the appliance. The table includes static and default routes from interface configuration and dynamic routes from discovery protocols.

GET

Use the HTTP GET method with the `/mgmt/status/default/RoutingStatus3` resource to retrieve routing table details.

- [Resource URL](#)
- [“Request headers” on page 1140](#)
- [Request body format](#)
- [“Security requirements” on page 1140](#)
- [Response status codes](#)
- [“Response headers” on page 1141](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/RoutingStatus3
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about the routing table. Each object contains the following fields:

IPType

String.

The version of the IP address: `ipv4` or `ipv6`.

Destination

String.

The IP subnet or aggregate that is the destination of the route.

PrefixLength

Integer.

The prefix length of the IP address that represents the configured netmask.

InterfaceType

String.

The type of device interface used for this route. For example, `Ethernet` indicates that the interface is on the base Ethernet MAC interface and `VLAN` indicates that the interface is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.

- **Ethernet** Indicates that it is on the base Ethernet MAC interface.
- **Tunnel** Indicates that it is on an encapsulation interface.
- **VLAN** Indicates that it is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.
- **Aggregate** Indicates that it is on an IEEE 802.3ad Link Aggregation composed of one of the base Ethernet MAC interfaces.
- **Other** Refers to network interfaces that you cannot configure.

MacInterface

String.

The configured name of the interface for the route.

GatewayIPType

String.

The IP address version that corresponds to the next hop gateway: `ipv4` or `ipv6`.

Gateway

String.

The next hop gateway for traffic that uses the route. For routes with a local route type, the value is `0.0.0.0` for IPv4 and `::` for IPv6.

Metric

Integer.

The metric for the route. When the same destination has multiple routes, the route with the lowest metric is used.

Examples

The following example requests information about the routing table on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/RoutingStatus3
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/RoutingStatus3"
    },
    "doc": {
      "href": "/mgmt/docs/status/RoutingStatus3"
    }
  },
  "RoutingStatus3": [
    {
      "IPType": "ipv4",
      "Destination": "0.0.0.0",
      "PrefixLength": 0,
      "InterfaceType": "Ethernet",
      "MacInterface": "mgt0",
      "GatewayIPType": "ipv4",
      "Gateway": "9.20.87.1",
      "Metric": 200
    },
    {
      "IPType": "ipv4",
      "Destination": "9.20.87.0",
      "PrefixLength": 24,
      "InterfaceType": "Ethernet",
      "MacInterface": "mgt0",
      "GatewayIPType": "ipv4",
      "Gateway": "0.0.0.0",
      "Metric": 0
    },
    {
      "IPType": "ipv4",
      "Destination": "172.21.0.0",
      "PrefixLength": 16,
      "InterfaceType": "Ethernet",
      "MacInterface": "eth20",
      "GatewayIPType": "ipv4",
      "Gateway": "0.0.0.0",
      "Metric": 0
    },
    {
      "IPType": "ipv4",
      "Destination": "192.168.251.0",
      "PrefixLength": 24,
      "InterfaceType": "Ethernet",
      "MacInterface": "eth13",
      "GatewayIPType": "ipv4",
      "Gateway": "0.0.0.0",
      "Metric": 0
    },
    {
      "IPType": "ipv4",
      "Destination": "192.168.253.0",
      "PrefixLength": 24,
      "InterfaceType": "Ethernet",
      "MacInterface": "eth21",
      "GatewayIPType": "ipv4",
      "Gateway": "0.0.0.0",
      "Metric": 0
    }
  ]
}
```

[/mgmt/status/default/ServicesMemoryStatus2](#)

You can use the HTTP GET method with the `/mgmt/status/default/ServicesMemoryStatus2` resource to view information about memory usage by services on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/ServicesMemoryStatus2` resource to view information about memory usage by services on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1143](#)
- [Request body format](#)
- [“Security requirements” on page 1143](#)
- [Response status codes](#)
- [“Response headers” on page 1144](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/ServicesMemoryStatus2
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about memory usage by services. Each object contains the following fields:

serviceClass

String.

The type of service being reported on.

serviceName

String.

The name of the service.

current

Integer.

Current memory in use in megabytes.

oneMinute

Integer.

The peak memory usage by the service over the previous minute.

fiveMinutes

Integer.

The peak memory usage by the service in the interval between 1 and 5 minutes ago.

tenMinutes

Integer.

The peak memory usage by the service in the interval between 5 and 10 minutes ago.

oneHour

Integer.

The peak memory usage by the service in the interval between 10 minutes and 1 hour ago.

twelveHours

Integer.

The peak memory usage by the service in the interval between 1 and 12 hours ago.

oneDay

Integer.

The peak memory usage by the service in the interval between 12 hours and 1 day ago.

lifetime

Integer.

The peak memory usage by the service over its lifetime.

Examples

The following example requests information about service memory usage on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/ServicesMemoryStatus2
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/ServicesMemoryStatus2"
    }
  }
}
```

```
    },
    "doc": {
      "href": "/mgmt/docs/status/ServicesMemoryStatus2"
    }
  },
  "result": "No status retrieved."
}
```

[/mgmt/status/default/ServicesStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/ServicesStatus` resource to view information about active services on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/ServicesStatus` resource to view information about active services on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1145](#)
- [Request body format](#)
- [“Security requirements” on page 1145](#)
- [Response status codes](#)
- [“Response headers” on page 1146](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/ServicesStatus
```

Where `port` is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about active services on the appliance. Each object contains the following fields:

LocalIP

String.

The IP address being used by this service.

LocalPort

Integer.

The port being used by this service.

ServiceClass

String.

Type of service.

ServiceName

String.

Name of the service.

Examples

The following example requests information about active services on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/ServicesStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/ServicesStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/ServicesStatus"
    }
  },
  "ServicesStatus": [
    {
      "LocalIP": "0.0.0.0",
      "LocalPort": 5554,
      "ServiceClass": "MultiProtocolGateway",
      "ServiceName": "rest-mgmt"
    },
    {
      "LocalIP": "0.0.0.0",
      "LocalPort": 9090,
      "ServiceClass": "WebGUI",
      "ServiceName": "web-mgmt"
    }
  ]
}
```

[/mgmt/status/default/SNMPStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/SNMPStatus` resource to view the SNMP status on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/SNMPStatus` resource to view the SNMP status on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1147](#)
- [Request body format](#)
- [“Security requirements” on page 1147](#)
- [Response status codes](#)
- [“Response headers” on page 1148](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/SNMPStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about SNMP status and contains the following fields:

EngineID

String.

The engineID for the local SNMPv3 engine.

EngineBoots

Integer.

The number of reboots for the local SNMPv3 engine.

Examples

The following example requests information about SNMP status on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/SNMPStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/SNMPStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/SNMPStatus"
    }
  },
  "SNMPStatus": {
    "EngineID": "0x8000395d03000bab8a56ee",
    "EngineBoots": 427
  }
}
```

[/mgmt/status/default/SSHTrustedHostStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/SSHTrustedHostStatus` resource to view the SSH known host table on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/SSHTrustedHostStatus` resource to view the SSH known host table on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1149](#)
- [Request body format](#)
- [“Security requirements” on page 1149](#)
- [Response status codes](#)
- [“Response headers” on page 1149](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/SSHTrustedHostStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about known SSH hosts. Each object contains the following fields:

Host

String.

The hostname of the peer.

Type

String.

The key type recorded for peer.

Examples

The following example requests information about known SSH hosts on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/SSHTrustedHostStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/SSHTrustedHostStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/SSHTrustedHostStatus"
    }
  },
  "SSHTrustedHostStatus": [
    {
      "Host": "9.20.32.227",
      "Type": "ssh-rsa"
    },
    {
      "Host": "9.20.32.231",
      "Type": "ecdsa-sha2-nistp256"
    },
    {
      "Host": "9.20.33.218",
      "Type": "ecdsa-sha2-nistp256"
    },
    {
      "Host": "9.20.33.238",
      "Type": "ecdsa-sha2-nistp256"
    },
    {
      "Host": "9.20.33.238",
      "Type": "ssh-rsa"
    },
    {
      "Host": "9.20.33.8",
      "Type": "ssh-rsa"
    },
    {
      "Host": "9.20.58.175",
      "Type": "ssh-rsa"
    },
    {
      "Host": "9.20.87.170",
      "Type": "ssh-rsa"
    }
  ]
}
```

[/mgmt/status/default/SystemCpuStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/SystemCpuStatus` resource to request information about CPU usage on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/SystemCpuStatus` resource to request information about CPU usage on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1151](#)
- [Request body format](#)
- [“Security requirements” on page 1151](#)
- [Response status codes](#)
- [“Response headers” on page 1151](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/SystemCpuStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information CPU usage and contains the following fields:

CpuUsage

Integer.

The instantaneous CPU usage as a percentage of the CPU load.

CpuLoadAvg1

String.

Shows the average CPU load over the last minute.

CpuLoadAvg5

String.

Shows the average CPU load over the last 5 minutes.

CpuLoadAvg15

String.

Shows the average CPU load over the last 15 minutes.

Examples

The following example requests information about CPU usage on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/SystemCpuStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/SystemCpuStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/SystemCpuStatus"
    }
  },
  "SystemCpuStatus": {
    "CpuUsage": 2,
    "CpuLoadAvg1": "0.03",
    "CpuLoadAvg5": "0.04",
    "CpuLoadAvg15": "0.04"
  }
}
```

[/mgmt/status/default/SystemMemoryStatus](#)

You can use the HTTP GET method with the `/mgmt/status/default/SystemMemoryStatus` resource to request information about memory usage on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/SystemMemoryStatus` resource to request information about memory usage on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1152](#)
- [Request body format](#)
- [“Security requirements” on page 1153](#)
- [Response status codes](#)
- [“Response headers” on page 1153](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/SystemMemoryStatus
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about memory usage and contains the following fields:

MemoryUsage

Integer.

The instantaneous memory usage as a percentage of the total memory.

TotalMemory

Integer.

The total memory of the system in megabytes. The total memory equals the amount of installed memory minus the amount of reserved memory.

UsedMemory

Integer.

The amount of memory that is currently in use. The used memory equals the amount of total memory minus the amount of free memory. The used memory does not include any hold memory.

FreeMemory

Integer.

The amount of memory that is currently not in use and is therefore available. The free memory value includes any hold memory that is not currently in use.

Examples

The following example requests information about memory usage on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/SystemMemoryStatus
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/SystemMemoryStatus"
    },
    "doc": {
      "href": "/mgmt/docs/status/SystemMemoryStatus"
    }
  },
  "SystemMemoryStatus": {
    "MemoryUsage": 3,
    "TotalMemory": 193638,
    "UsedMemory": 4955,
    "FreeMemory": 188683
  }
}
```

[/mgmt/status/default/TCPSummary](#)

You can use the HTTP GET method with the `/mgmt/status/default/TCPSummary` resource to list the number of TCP connections on this appliance in each of the states defined in the TCP specification, RFC 793.

GET

Use the HTTP GET method with the `/mgmt/status/default/TCPSummary` resource to list the number of TCP connections on this appliance in each of the states defined in the TCP specification, RFC 793.

- [Resource URL](#)
- [“Request headers” on page 1154](#)
- [Request body format](#)
- [“Security requirements” on page 1154](#)
- [Response status codes](#)
- [“Response headers” on page 1155](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/TCPSummary
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about TCP connections and contains the following fields:

established

Integer.

The number of TCP connections in the established state. Connections in this state have completed all handshakes and can transfer data in either direction.

syn_sent

Integer.

The number of TCP connections in the syn-sent state. Connections in this state are waiting for a matching connection request after sending a connection request.

syn_received

Integer.

The number of TCP connections in the syn-received state. Connections in this state are waiting for a confirming connection request acknowledgment after both receiving and sending a connection request.

fin_wait_1

Integer.

The number of TCP connections in the fin-wait-1 state. Connections in this state are waiting for a connection termination request from the remote TCP or an acknowledgment of the connection termination request previously sent.

fin_wait_2

Integer.

The number of TCP connections in the fin-wait-2 state. Connections in this state are waiting for a connection termination request from the remote TCP.

time_wait

Integer.

The number of TCP connections in the time-wait state. Connections in this state are waiting for enough time to pass to be sure that the remote TCP received the acknowledgment of its connection termination request.

closed

Integer.

The number of TCP connections in the closed state. This state represents no connection state at all.

close_wait

Integer.

The number of TCP connections in the close-wait state. Connections in this state are waiting for a connection termination request from the local user.

last_ack

Integer.

The number of TCP connections in the last-ack state. Connections in this state are waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request).

listen

Integer.

The number of TCP connections in the listen state. Connections in the listen state are waiting for a connection request from any remote TCP and port.

closing

Integer.

The number of TCP connections in the closing state. Connections in this state are waiting for a connection termination request acknowledgment from the remote TCP.

Examples

The following example requests information about TCP connections on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/TCPSummary
```

The following listing shows an example received response:

```

{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/TCPSummary"
    },
    "doc": {
      "href": "/mgmt/docs/status/TCPSummary"
    }
  },
  "TCPSummary": {
    "established": 3,
    "syn_sent": 0,
    "syn_received": 0,
    "fin_wait_1": 0,
    "fin_wait_2": 0,
    "time_wait": 0,
    "closed": 0,
    "close_wait": 0,
    "last_ack": 0,
    "listen": 7,
    "closing": 0
  }
}

```

[/mgmt/status/default/TCPTable](#)

You can use the HTTP GET method with the `/mgmt/status/default/TCPTable` resource to list all TCP connections on the system. Active connections and listeners are included. For listeners that are related to router configuration objects, the table indicates with what object the listener is associated.

GET

Use the HTTP GET method with the `/mgmt/status/default/TCPTable` resource to list all TCP connections on the system.

- [Resource URL](#)
- [“Request headers” on page 1157](#)
- [Request body format](#)
- [“Security requirements” on page 1157](#)
- [Response status codes](#)
- [“Response headers” on page 1157](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/TCPTable
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about TCP connections. Each object contains the following fields:

localIP

String.

The local IP address of this TCP connection. For connections in the listen state, an IP address of 0.0.0.0 means that connections will be accepted at any IP address of the system. A specific IP address indicates that connections will only be accepted at that IP address.

localPort

Integer.

The local TCP port number of this TCP connection.

remoteIP

String.

The remote IP address of this TCP connection. The value is 0.0.0.0 for connections in the listen state. The value is a specific IP address for established connections.

remotePort

Integer.

The remote TCP port number of this TCP connection.

state

String.

The state of the TCP connection as defined in RFC 793:

- established
- syn-sent
- syn-received
- fin-wait-1
- fin-wait-2
- time-wait
- closed
- close-wait
- last-ack
- listen
- closing

serviceDomain

String.

For TCP connections that were created by a configuration object this value indicates the domain that the configuration object that created the listener is in. Only applies to connections in the listen state.

serviceClass

String.

The type of configuration object that created the listener. For all connections not in the listen state this value will be configBase.

serviceName

String.

The name of the configuration object that created the listener. Only applies to TCP connections in the listen state.

Examples

The following example requests information about TCP connections on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/TCPTable
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/TCPTable"
    },
    "doc": {
      "href": "/mgmt/docs/status/TCPTable"
    }
  },
  "TCPTable": [
    {
      "localIP": "0.0.0.0",
      "localPort": 22,
      "remoteIP": "0.0.0.0",
      "remotePort": 0,
      "state": "listen",
      "serviceDomain": "default",
      "serviceClass": "SSHService",
      "serviceName": "SSH Service"
    },
    {
      "localIP": "0.0.0.0",
      "localPort": 5554,
      "remoteIP": "0.0.0.0",
      "remotePort": 0,
      "state": "listen",
      "serviceDomain": "default",
      "serviceClass": "HTTPSSourceProtocolHandler",
      "serviceName": "rest-mgmt-https"
    },
    {
      "localIP": "0.0.0.0",
      "localPort": 9090,
      "remoteIP": "0.0.0.0",
      "remotePort": 0,
      "state": "listen",
      "serviceDomain": "default",
      "serviceClass": "WebGUI",
      "serviceName": "web-mgmt"
    },
    {
      "localIP": "9.20.87.136",
      "localPort": 5554,
      "remoteIP": "9.140.111.179",
      "remotePort": 55888,
      "state": "established",
      "serviceDomain": "",
      "serviceClass": "None",
      "serviceName": ""
    },
    {
      "localIP": "127.0.0.1",
      "localPort": 6666,
      "remoteIP": "0.0.0.0",
      "remotePort": 0,
      "state": "listen",
      "serviceDomain": "",
      "serviceClass": "None",
      "serviceName": ""
    },
    {
      "localIP": "::",
      "localPort": 26360,
      "remoteIP": "::",
      "remotePort": 0,
      "state": "listen",
      "serviceDomain": "",
      "serviceClass": "None",
      "serviceName": ""
    },
    {
      "localIP": "127.0.0.1",
      "localPort": 9443,
      "remoteIP": "::",

```

```

    "remotePort": 0,
    "state": "listen",
    "serviceDomain": "",
    "serviceClass": "None",
    "serviceName": ""
  },
  {
    "localIP": "127.0.0.1",
    "localPort": 9444,
    "remoteIP": "::",
    "remotePort": 0,
    "state": "listen",
    "serviceDomain": "",
    "serviceClass": "None",
    "serviceName": ""
  },
  {
    "localIP": "127.0.0.1",
    "localPort": 58606,
    "remoteIP": "127.0.0.1",
    "remotePort": 63503,
    "state": "established",
    "serviceDomain": "",
    "serviceClass": "None",
    "serviceName": ""
  }
]
}

```

[/mgmt/status/default/TemperatureSensors](#)

You can use the HTTP GET method with the `/mgmt/status/default/TemperatureSensors` resource to get information from sensors that read temperatures, both of the air flowing through the appliance and of key components in the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/TemperatureSensors` resource to get information from sensors that read temperatures.

- [Resource URL](#)
- [“Request headers” on page 1160](#)
- [Request body format](#)
- [“Security requirements” on page 1161](#)
- [Response status codes](#)
- [“Response headers” on page 1161](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/TemperatureSensors
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about temperature sensors. Each object contains the following fields:

Name

String.

The name of the temperature sensor that is being monitored.

Value

Integer.

The most recent reading of the temperature sensor in degree Celsius. There are only three significant digits.

UpperNonCriticalThreshold

Integer.

The temperature at which a warning of high temperature occurs. If the current temperature is above this value, investigate the cause and correct the problem.

UpperNonRecoverableThreshold

Integer.

The temperature at which a risk of permanent damage to the appliance exists. If the current temperature is above this threshold, correct the problem immediately; otherwise, turn off the power supply of the appliance.

ReadingStatus

String.

An indicator of whether the current temperature is OK or exceeds the warning, danger, or critical threshold.

lowerNonRecoverable

The current sensor reading value is below the lower non-recoverable threshold, which indicates that there is a risk of damaging the hardware if the device is allowed to remain in this state. If you cannot diagnose and fix the cause of this reading, you should power down the device.

lowerCritical

The current sensor reading value is below the lower critical threshold, which indicates that the system is faulty, and might not be operating properly.

lowerNonCritical

The current sensor reading value is below the lower non-critical threshold, which is a warning that it is getting close to too low. The system should operate normally, but the cause of the sensor value should be investigated.

ok

The current sensor reading value is OK, and is within all limits.

upperNonCritical

The current sensor reading value is above the upper non-critical threshold, which is a warning that it is getting close to too high. The system should operate normally, but the cause of the sensor value should be investigated. This is commonly used on temperature sensors, where the investigation would be checking ambient air temperature and air flow.

upperCritical

The current sensor reading value is above the upper critical threshold, which indicates that the system is faulty, and might not be operating properly.

upperNonRecoverable

The current sensor reading value is above the upper non-recoverable threshold, which indicates that there is a risk of damaging the hardware if the device is allowed to remain in this state. If you cannot diagnose and fix the cause of this reading, you should power down the device.

failure

The reading indicates a failure. Used for non-numeric readings.

noReading

The attempt to read the sensor failed, so the status cannot be checked.

invalid

An internal software error has been detected checking the status.

UpperCriticalThreshold

Integer.

The temperature at which a critical error of high temperature occurs. If the current temperature is above this value, correct the problem immediately.

Examples

The following example requests information about temperature sensors on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/TemperatureSensors
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/TemperatureSensors"
    },
    "doc": {
      "href": "/mgmt/docs/status/TemperatureSensors"
    }
  },
  "TemperatureSensors": [
    {
      "Name": "CPU 1 DIMM Channel 0 Temperature",
      "Value": 34,
```

```

    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "CPU 1 DIMM Channel 1 Temperature",
    "Value": 34,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "CPU 1 DIMM Channel 2 Temperature",
    "Value": 34,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "CPU 1 DIMM Channel 3 Temperature",
    "Value": 33,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "CPU 2 DIMM Channel 0 Temperature",
    "Value": 33,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "CPU 2 DIMM Channel 1 Temperature",
    "Value": 33,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "CPU 2 DIMM Channel 2 Temperature",
    "Value": 33,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "CPU 2 DIMM Channel 3 Temperature",
    "Value": 32,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "Temperature CPU1",
    "Value": 34,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "Temperature CPU2",
    "Value": 37,
    "UpperNonCriticalThreshold": 85,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 90
  },
  {
    "Name": "Temperature System 1",
    "Value": 30,
    "UpperNonCriticalThreshold": 55,
    "UpperNonRecoverableThreshold": "",

```

```

    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 65
  },
  {
    "Name": "Temperature System 2",
    "Value": 34,
    "UpperNonCriticalThreshold": 65,
    "UpperNonRecoverableThreshold": "",
    "ReadingStatus": "ok",
    "UpperCriticalThreshold": 75
  }
]
}

```

[/mgmt/status/default/VlanInterfaceStatus2](#)

You can use the HTTP GET method with the `/mgmt/status/default/VlanInterfaceStatus2` resource to retrieve status for VLAN interfaces on the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/VlanInterfaceStatus2` resource to retrieve status for VLAN interfaces on the appliance.

- [Resource URL](#)
- [“Request headers” on page 1164](#)
- [Request body format](#)
- [“Security requirements” on page 1164](#)
- [Response status codes](#)
- [“Response headers” on page 1165](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/VlanInterfaceStatus2
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated”](#) on page 423 for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles”](#) on page 423 for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about VLAN interfaces. Each object contains the following fields:

Name

String.

The name of the interface in the system kernel.

MgmtName

String.

The name of the interface in the configuration.

InterfaceIndex

Integer.

The SNMP ifIndex value of the row in the SNMP ifTable for the VLAN interface.

OnIfIndex

Integer.

The SNMP ifIndex of the parent interface of the VLAN interface.

OnIfName

String.

The name of the parent interface of the VLAN interface.

VlanIdentifier

String.

The VLAN identifier that uniquely identifies the VLAN interface on its associated physical Ethernet network.

Examples

The following example requests information about VLAN interfaces on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/VlanInterfaceStatus2
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/VlanInterfaceStatus2"
    },
    "doc": {
      "href": "/mgmt/docs/status/VlanInterfaceStatus2"
    }
  },
}
```

```
"result": "No status retrieved."  
}
```

[/mgmt/status/default/VoltageSensors](#)

You can use the HTTP GET method with the `/mgmt/status/default/VoltageSensors` resource to get information from sensors that read voltages, from the power supplies and other components of the appliance.

GET

Use the HTTP GET method with the `/mgmt/status/default/VoltageSensors` resource to get information from sensors that read voltages, from the power supplies and other components of the appliance.

- [Resource URL](#)
- [“Request headers” on page 1166](#)
- [Request body format](#)
- [“Security requirements” on page 1166](#)
- [Response status codes](#)
- [“Response headers” on page 1167](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/status/default/VoltageSensors
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Status information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not have the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes information about voltage sensors. Each object contains the following fields:

Name

String.

The name of the voltage sensor that is being monitored.

Value

Integer.

The most recent reading of the voltage sensor in millivolts. There are only three significant digits.

LowerCriticalThreshold

Integer.

The lowest allowable reading of the voltage sensor in millivolts.

UpperCriticalThreshold

Integer.

The highest allowable reading of the voltage sensor in millivolts.

ReadingStatus

String.

An indicator of whether the current voltage is OK or the voltage exceeds the upper or lower critical threshold.

lowerNonRecoverable

The current sensor reading value is below the lower non-recoverable threshold, which indicates that there is a risk of damaging the hardware if the device is allowed to remain in this state. If you cannot diagnose and fix the cause of this reading, you should power down the device.

lowerCritical

The current sensor reading value is below the lower critical threshold, which indicates that the system is faulty, and might not be operating properly.

lowerNonCritical

The current sensor reading value is below the lower non-critical threshold, which is a warning that it is getting close to too low. The system should operate normally, but the cause of the sensor value should be investigated.

ok

The current sensor reading value is OK, and is within all limits.

upperNonCritical

The current sensor reading value is above the upper non-critical threshold, which is a warning that it is getting close to too high. The system should operate normally, but the cause of the sensor value should be investigated. This is commonly used on temperature sensors, where the investigation would be checking ambient air temperature and air flow.

upperCritical

The current sensor reading value is above the upper critical threshold, which indicates that the system is faulty, and might not be operating properly.

upperNonRecoverable

The current sensor reading value is above the upper non-recoverable threshold, which indicates that there is a risk of damaging the hardware if the device is allowed to remain in this state. If you cannot diagnose and fix the cause of this reading, you should power down the device.

failure

The reading indicates a failure. Used for non-numeric readings.

noReading

The attempt to read the sensor failed, so the status cannot be checked.

invalid

An internal software error has been detected checking the status.

Examples

The following example requests information about voltage sensors on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/status/default/VoltageSensors
```

The following listing shows an example received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/status/default/VoltageSensors"
    },
    "doc": {
      "href": "/mgmt/docs/status/VoltageSensors"
    }
  },
  "VoltageSensors": [
    {
      "Name": "Voltage +12",
      "Value": 11900,
      "LowerCriticalThreshold": 10800,
      "UpperCriticalThreshold": 13200,
      "ReadingStatus": "ok"
    },
    {
      "Name": "Voltage +3.3",
      "Value": 3310,
      "LowerCriticalThreshold": 2980,
      "UpperCriticalThreshold": 3650,
      "ReadingStatus": "ok"
    },
    {
      "Name": "Voltage +5",
      "Value": 5060,
      "LowerCriticalThreshold": 4480,
      "UpperCriticalThreshold": 5500,
      "ReadingStatus": "ok"
    },
    {
      "Name": "Voltage +5 Standby",
      "Value": 5060,
      "LowerCriticalThreshold": 4480,
      "UpperCriticalThreshold": 5500,
      "ReadingStatus": "ok"
    },
    {
      "Name": "Voltage Battery",
      "Value": 3120,
      "LowerCriticalThreshold": 2690,
      "UpperCriticalThreshold": 3460,
      "ReadingStatus": "ok"
    },
    {
      "Name": "Voltage CPU Bus Termination",
      "Value": 1010,
      "LowerCriticalThreshold": 944,
      "UpperCriticalThreshold": 1150,
      "ReadingStatus": "ok"
    }
  ]
}
```

```

{
  "Name": "Voltage CPU1 Core",
  "Value": 816,
  "LowerCriticalThreshold": 496,
  "UpperCriticalThreshold": 1340,
  "ReadingStatus": "ok"
},
{
  "Name": "Voltage CPU2 Core",
  "Value": 816,
  "LowerCriticalThreshold": 496,
  "UpperCriticalThreshold": 1340,
  "ReadingStatus": "ok"
},
{
  "Name": "Voltage Memory Bus Termination",
  "Value": 1010,
  "LowerCriticalThreshold": 944,
  "UpperCriticalThreshold": 1150,
  "ReadingStatus": "ok"
},
{
  "Name": "Voltage PU +12",
  "Value": 12200,
  "LowerCriticalThreshold": 11400,
  "UpperCriticalThreshold": 13000,
  "ReadingStatus": "ok"
},
{
  "Name": "Voltage PU +3.3",
  "Value": 3400,
  "LowerCriticalThreshold": 2800,
  "UpperCriticalThreshold": 3600,
  "ReadingStatus": "ok"
},
{
  "Name": "Voltage PU +5",
  "Value": 5200,
  "LowerCriticalThreshold": 4600,
  "UpperCriticalThreshold": 5440,
  "ReadingStatus": "ok"
}
]
}

```

Actionqueue resources

The following URIs give access to the actionqueue resources on the appliance.

The actionqueue resources are used to trigger operations on the appliance.

Related concepts

[“Triggering appliance operations by using the REST management interface” on page 370](#)

A subset of appliance operations can be triggered by using the REST management interface.

[/mgmt/actionqueue/default/operations](#)

You can use the HTTP GET method with the `/mgmt/actionqueue/default/operations` resource to list all available operations that you can trigger on the appliance.

GET

Use the HTTP GET method with the `/mgmt/actionqueue/default/operations` resource to request a list of appliance operations.

- [Resource URL](#)
- [“Request headers” on page 1170](#)
- [Request body format](#)
- [“Security requirements” on page 1170](#)
- [Response status codes](#)
- [“Response headers” on page 1170](#)
- [Response body format](#)

- [Examples](#)

Resource URL

```
https://host:port/mgmt/actionqueue/default/operations
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Information retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response includes a list of the operations that you can trigger by using the `actionqueue` resource.

Examples

The following example requests a list of the operations that you can trigger on the appliance.

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/actionqueue/default/operations
```

The following listing shows an extract from the received response:

```

{
  "_links": {
    "self": {
      "href": "/mgmt/actionqueue/default/operations"
    },
    "AddPasswordMap": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/AddPasswordMap"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/AddPasswordMap"
      }
    },
    "AddSelTestEntry": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/AddSelTestEntry"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/AddSelTestEntry"
      }
    },
    "AddTrustedHost": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/AddTrustedHost"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/AddTrustedHost"
      }
    },
    "ApplyPatch": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/ApplyPatch"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/ApplyPatch"
      }
    },
    "BootDelete": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/BootDelete"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/BootDelete"
      }
    },
    "BootSwitch": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/BootSwitch"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/BootSwitch"
      }
    },
    "ChangePassword": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/ChangePassword"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/ChangePassword"
      }
    },
    "ClearSel": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/ClearSel"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/ClearSel"
      }
    },
    "ConvertCertificate": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/ConvertCertificate"
      },
      "metadata": {
        "href": "/mgmt/metadata/default/operations/ConvertCertificate"
      }
    },
    "ConvertKey": {
      "schema": {
        "href": "/mgmt/actionqueue/default/operations/ConvertKey"
      }
    }
  }
}

```


- [“Response headers” on page 1173](#)
- [Response body format](#)
- [Examples](#)

Resource URL

```
https://host:port/mgmt/metadata/default/operations/{operation}
```

Where *port* is 5554 by default.

Request headers

The following header must be sent with the request:

Authorization

This header must be sent to perform authentication.

Request body format

None.

Security requirements

The caller must be authenticated as an appliance user with sufficient authority. For more information about security, see [“Types of user and how they are authenticated” on page 423](#) and [“User authorization, credential mapping, and access profiles” on page 423](#).

Response status codes

200

Metadata retrieved successfully.

400

Invalid data provided.

401

Not authenticated.

The caller must be authenticated as an appliance user. See [“Types of user and how they are authenticated” on page 423](#) for more information.

403

Not authorized.

The caller is authenticated but does not the authority to perform this action. See [“User authorization, credential mapping, and access profiles” on page 423](#) for more information.

Response headers

Content-Type

This header is returned with a value of `application/json`.

Response body format

The response is in JSON format. The response gives the metadata for the specified operation.

The JSON object named `action` contains the following attributes which you need to form the POST request.

name

String.

Specifies the name of the operation.

parameters

Specifies any parameters that the operation requires.

Examples

The following example requests the metadata for the AddPasswordMap operation:

The following URL is used with the HTTP GET method:

```
https://localhost:5554/mgmt/metadata/default/operations/AddPasswordMap
```

The following listing shows the received response:

```
{
  "_links": {
    "self": {
      "href": "/mgmt/metadata/default/operations/AddPasswordMap"
    },
    "doc": {
      "href": "/mgmt/docs/metadata/AddPasswordMap"
    }
  },
  "action": {
    "name": "AddPasswordMap",
    "uri": "crypto/add-password-map",
    "cmd-group": "crypto",
    "cli-alias": "add password-map",
    "parameters": {
      "parameter": [
        {
          "name": "AliasName",
          "required": "true",
          "type": "dmString"
        },
        {
          "name": "Password",
          "required": "true",
          "type": "dmString"
        }
      ]
    },
    "display": "Add Password Map",
    "summary": "Add new password to password map",
    "description": "Adds a password to the existing password map. The password alias can be used in the configuration to refer to the password."
  }
}
```

Related concepts

[“Triggering appliance operations by using the REST management interface” on page 370](#)

A subset of appliance operations can be triggered by using the REST management interface.

[/mgmt/actionqueue/default](#)

You can use the HTTP POST method with the `/mgmt/actionqueue/default` resource to trigger an appliance action and specify any required parameters.

You specify the action and the parameters in the request body.

Related concepts

[“Triggering appliance operations by using the REST management interface” on page 370](#)

A subset of appliance operations can be triggered by using the REST management interface.

POST

Use the HTTP POST method with the `actionqueue` resource to trigger appliance operations.

The operation to be triggered is specified in the request body.

- [Resource URL](#)
- [Optional query parameters](#)
- [“Request headers” on page 1175](#)

- [Request body format](#)
- [Response status codes](#)
- [“Response headers” on page 1175](#)
- [Response body format](#)
- [Examples](#)

Resource URL

`https://host:port/mgmt/actionqueue/default/`

Where *port* is 5554 by default.

Optional query parameters

None.

Request headers

The following headers must be sent with the request:

Content-Type

This header must be sent with a value of `application/json`.

Request body format

The request body must be in JSON format in UTF-8 encoding. The response body specifies the operation and any required parameters. See the operation-specific subtopics for details.

You can view information about the various data types expected in the POST bodies by using the `/mgmt/types/resource` (see [“REST management resources summary” on page 1002](#)).

Response status codes

200

Operation succeeded.

400

Invalid data provided.

401

Not authenticated.

An invalid user name or password was provided.

500

Server issue.

Response headers

None.

Response body format

The response body contains information generated by the action.

Examples

Examples are given in the operation-specific subtopics.

AddPasswordMap

Trigger this action to add a password to the password map.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "AddPasswordMap": {
    "AliasName": "password_alias",
    "Password": "password"
  }
}
```

AliasName

String

The alias is a publicly known reference that is included in configuration files and exports. With an alias you keep the real password secret because the plain text value for the password is stored in an encrypted file. The password for the alias cannot be viewed or accessed by anyone. Only the appliance can extract the password that it uses internally.

Password

String

The plain text password to alias.

Example

To set the alias "wills_password" for the plain text password "abc_123_def_456", post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "AddPasswordMap": {
    "AliasName": "wills_password",
    "Password": "abc_123_def_456"
  }
}
```

AddSelTestEntry

Trigger this action to add an event to the IPMI SEL log for test purposes.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "AddSelTestEntry": {
  }
}
```

This action adds an event to the Intelligent Platform Management Interface (IPMI) System Event Log (SEL). SEL events are collected by the Baseboard Management Controller (BMC). Use this action to test the code that captures events from the SEL. When run, this action creates an OEM SEL record of hex type 'DF' with an OEM-defined value of 'DPTTest'.

AddTrustedHost

Trigger this action to define an SSH peer as an SSH known host.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "AddTrustedHost": {
    "Host": "IP_address",
    "Type": "key_type",
    "Key": "key"
  }
}
```

Host

String

The IP address for the SSH peer.

Type

String

The key type for peer.

Key

String

The host key value for the SSH peer.

Example

To set the host with the IP address 10.97.111.108 as a trusted host, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "AddTrustedHost": {
    "Host": "10.97.111.108",
    "Type": "ssh-rsa",
    "Key": "AAAAB3NzaC1yc2EAAAABIwAAIEA1J/99rRvdZmVvkaKvcG2a+PeCm25
p80Jl87SA6mtFxudA2ME6n3lcXEakpQ8KFTpPbBXt+yDKNFR9gNHIfR1
UDho1HAN/a0gEsvrnDY5wKrTcRHrqDc/x0buPzbsEmXi01ud5P17+BXQ
VpPbyVujoHINCrx0k/z7Qpkozb4qZd8=="
  }
}
```

ApplyPatch

Trigger this action to designate the startup firmware image and restart the appliance with this image.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "ApplyPatch": {
    "AcceptLicense": "accept-license",
    "File": "file_name"
  }
}
```

AcceptLicense

String

Specify "accept-license" to accept the license.

File

String

Specifies the name of the firmware image to boot the appliance with.

Example

To specify the image in the file 180858 and restart the appliance, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "ApplyPatch": {
    "AcceptLicense": "accept-license",
    "File": "180858"
  }
}
```

BootDelete

Trigger this action to delete the secondary installation. Do not delete the secondary installation unless directed by IBM® Support.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "BootDelete": {
  }
}
```

A firmware upgrade with the ApplyPatch action 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.

BootSwitch

Trigger this action to switch between primary and secondary installations.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "BootSwitch": {
  }
}
```

A firmware upgrade with the ApplyPatch action 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.

The boot switch action is not available on an appliance that is part of a high availability group.

ChangePassword

Trigger this action to change your password on the appliance.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "ChangePassword": {
    "OldPassword": "old_password",
    "Password": "new_password"
  }
}
```

OldPassword

String

The existing password.

Password

String

The new password.

Example

To change your password on the appliance from "abc_123_def_345" to "xyz_99_rst_00", post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{ "ChangePassword": {
  "OldPassword": "abc_123_def_345",
  "Password": "xyz_99_rst_00"
}
}
```

ClearSel

Trigger this action to clear the IPMI SEL log.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "ClearSel": {
  }
}
```

This action deletes all events from the Intelligent Platform Management Interface (IPMI) System Event Log (SEL). SEL events are collected by the Baseboard Management Controller (BMC).

ConvertCertificate

Trigger this action to write a Crypto Certificate object to a file in a specific format.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "ConvertCertificate": {
    "ObjectName": "object",
    "OutputFilename": "file",
    "Format": "format"
  }
}
```

ObjectName

String

Specifies the name of the certificate alias.

OutputFilename

String

Specifies the output file name. Use the `temporary:///mycert.pub` format.

Format

String

Specifies the format for the output file. The supported format is `openssh-pubkey`.

This action 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.

Example

To save the certificate named `iop_mgmt_cert` to the file `savedthecert.pub`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "ConvertCertificate": {
    "ObjectName": "iop_mgmt_cert",
    "OutputFilename": "temporary:///savedthecert.pub",
    "Format": "openssh-pubkey"
  }
}
```

ConvertKey

Trigger this action to write a Crypto Key object to a file in a specific format.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "ConvertKey": {
    "ObjectName": "object",
    "OutputFilename": "file",
  }
}
```

```
    "Format": "format"
  }
}
```

ObjectName

String

Specifies the name of the key alias.

OutputFilename

String

Specifies the output file name. Use the `temporary:///mykey.pub` format.

Format

String

Specifies the format for the output file. The supported format is `openssh-pubkey`.

This action 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.

Example

To save the key named `iop_mgmt_key` to the file `savedthekey.pub`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "ConvertKey": {
    "ObjectName": "iop_mgmt_key",
    "OutputFilename": "temporary:///savedthekey.pub",
    "Format": "openssh-pubkey"
  }
}
```

CreateDir

Trigger this action to create a subdirectory in the `local:///` directory.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "CreateDir": {
    "Dir": "subdirectory"
  }
}
```

Dir

String

Subdirectory to create in the form `local:///subdirectory`.

Example

To create a directory named `mydir`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "CreateDir": {
    "Dir": "local:///mydir"
  }
}
```

CryptoExport

Trigger this action to create an export package that contains a certificate.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "CryptoExport": {
    "ObjectType": "type",
    "ObjectName": "object",
    "OutputFilename": "file",
    "Mechanism": "format"
  }
}
```

ObjectType

String

Always "certificate".

ObjectName

String

The name of the certificate to export.

OutputFilename

String

Specify the name of the export package. This file will contain the exported cryptographic object.

Example

To export a certificate named `iop_mgmt_cert` to the file `exportedthecert.xml`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "CryptoExport": {
    "ObjectType": "certificate",
    "ObjectName": "iop_mgmt_cert",
    "OutputFilename": "temporary:///exportedthecert.xml"
  }
}
```

CryptoImport

Trigger this action to import an export package that contains a certificate.

Use the HTTP POST method with the resource `/mgmt/metadata/default`, specifying a request payload with the following format:

```
{
  "CryptoImport": {
    "ObjectType": "type",
    "ObjectName": "object",
    "InputFilename": "file",
    "ImportPassword": "password",
    "ImportPasswordAlias": "password_alias"
  }
}
```

ObjectType

String

Always "certificate".

ObjectName

String

The name of the certificate to import.

InputFilename

String

Specify the name of the file to import. This file contains the cryptographic object to be imported.

ImportPassword

String

Optional: Specifies the password that was used to encrypt the input file. This parameter is mutually exclusive with the `ImportPasswordAlias` parameter.

ImportPasswordAlias

String

Optional: Specifies the password that was used to encrypt the input file. This parameter is mutually exclusive with the `ImportPassword` parameter.

Example

To import a certificate named `iop_mgmt_cert` from the file `exportedthecert.xml`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "CryptoImport": {
    "ObjectType": "certificate",
    "ObjectName": "iop_mgmt_cert",
    "InputFilename": "temporary:///exportedthecert.xml",
    "ImportPassword": "abc_123_def_456"
  }
}
```

CryptoModeSet

Trigger this action to set the appliance-wide cryptographic mode for the next firmware reload.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "CryptoModeSet": {
    "Mode": "crypto-mode"
  }
}
```

Mode

String

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-140-2-11` - Runs the firmware in FIPS 140-2 Level 1 mode.

Example

To set the appliance cryptographic mode to the FIPS 140-2 Level 1 mode, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "CryptoModeSet": {
    "Mode": "fips-140-2-11"
  }
}
```

DeleteFile

Trigger this action to delete a file from the appliance.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "DeleteFile": {
    "File": "filepath"
  }
}
```

File

String

Specifies the location as a URL of the file to delete in the `directory:///file` format.

Example

To delete the file named `badfile` from the `temporary:///` directory, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "DeleteFile": {
    "File": "temporary:///badfile"
  }
}
```

DeletePasswordMap

Trigger this action to remove an entry from the password map.

Use the HTTP POST method with the resource `/mgmt/metadata/default`, specifying a request payload with the following format:

```
{
  "DeletePasswordMap": {
    "AliasName": "filepath"
  }
}
```

AliasName

String

Identifies the entry to delete from the password map.

Example

To delete the password alias named `"wills_password"`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "DeletePasswordMap": {
    "AliasName": "wills_password"
  }
}
```

DeleteTrustedHost

Trigger this action to remove an SSH peer as an SSH known host.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "DeleteTrustedHost": {
    "Host": "host"
  }
}
```

Host

String

Specifies the fully-qualified host name or IP address for the peer.

Example

To delete the host with the IP address 10.97.111.108 as a trusted host, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "DeleteTrustedHost": {
    "Host": "10.97.111.108",
  }
}
```

DisableEthernetHardwareOffload

Trigger this action to temporarily disable Ethernet hardware offload for the specified interface.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "DisableEthernetHardwareOffload": {
    "Interface": "Ethernet_interface"
  }
}
```

Interface

String

Specifies the Ethernet interface.

Example

To disable Ethernet hardware offload for the interface Eth13, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "DisableEthernetHardwareOffload": {
    "Interface": "eth13"
  }
}
```

DisableLinkAggregationHardwareOffload

Trigger this action to temporarily disable Link Aggregation hardware offload for the specified interface.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "DisableLinkAggregationHardwareOffload": {
    "Interface": "LinkAggregation_interface"
  }
}
```

Interface

String

Specifies the Link Aggregation interface.

Example

To disable hardware offload for the link aggregation interface LA203, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "DisableLinkAggregationHardwareOffload": {
    "Interface": "LA203"
  }
}
```

DisableVLANHardwareOffload

Trigger this action to temporarily disable VLAN hardware offload for the specified interface.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "DisableVLANHardwareOffload": {
    "Interface": "VLAN_interface"
  }
}
```

Interface

String

Specifies the VLAN interface.

Example

To disable hardware offload for the VLAN interface VLAN09, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "DisableVLANHardwareOffload": {
    "Interface": "VLAN09"
  }
}
```

Disconnect

Trigger this action to close the session for an active user.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "Disconnect": {
    "id": "sessionID", }
}
```

id

String

Specifies the session id of the session to end.

Use the `ActiveUsers` status resource to obtain a list of users currently logged in, together with session id, see ["/mgmt/status/default/ActiveUsers" on page 1008](#).

Example

To disconnect the session identified by session id 23, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "Disconnect": {
    "id": "23"
  }
}
```

```
}  
}
```

ErrorReport

Trigger this action to create an error report in the temporary: `///error-report.txt` file.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{  
  "ErrorReport": {}  
}
```

This action automatically creates a file that contains any backtrace or watchdog error report, the contents of the audit log, and the running configuration.

ExecConfig

Trigger this action to execute a configuration file.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{  
  "ExecConfig": {  
    "URL": "filepath"  
  }  
}
```

URL

String

Specifies the location of the configuration file as a URL in the *directory:///file* format.

Example

To execute the configuration file `config:///housekeeping.cfg`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{  
  "ExecConfig": {  
    "URL": "config:///housekeeping.cfg"  
  }  
}
```

FlushArpCache

Trigger this action to flush the ARP cache.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{  
  "FlushArpCache": {}  
}
```

FlushDNSCache

Trigger this action to flush the DNS cache.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{  
  "FlushDNSCache": {}  
}
```

FlushRBMCache

Trigger this action to flush the RBM cache.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "FlushRBMCache": {}
}
```

Keygen

Trigger this action to generate a public-private key pair and Certificate Signing Request (CSR). It can also be used to generate a self-signed certificate.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "Keygen": {
    "LDAPOrder": "yes" | "no",
    "C": "countryName",
    "ST": "stateOrProvinceName",
    "L": "localityName",
    "O": "organizationName",
    "OU": "organizationalUnitName",
    "OU1": "organizationalUnitName",
    "OU2": "organizationalUnitName",
    "OU3": "organizationalUnitName",
    "CN": "commonName",
    "KeyType": "keyType",
    "KeyLength": "keyLength",
    "Digest": "hashAlgorithm",
    "ECDSACurve": "ellipticCurves",
    "FileName": "fileName",
    "Days": "numberOfDays",
    "PasswordAlias": "alias",
    "ExportKey": "yes" | "no",
    "GenSSCert": "yes" | "no",
    "ExportSSCert": "yes" | "no",
    "GenObject": "yes" | "no",
    "ObjectName": "yes" | "no",
    "HSM": "yes" | "no",
    "UsingKey": "name",
    "ConfigFile": "fileName"
  }
}
```

LDAPOrder

Enum

When set to "on" specifies that the LDAP entry is created in reverse RDN order. The default setting is "off".

C

String

Optionally specifies the ISO two-character country identifier for the CSR.

ST

String

Optionally specifies the unabbreviated state or province name for the CSR. Use a text string up to 64 characters in length.

L

String

Optionally specifies the city or town name for the CSR. Use a text string up to 64 characters in length.

O

String

Optionally specifies the organization name for the CSR. Use a text string up to 64 characters in length.

OU

String

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.

OU1

String

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.

OU2

String

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.

OU3

String

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

String

Specifies the fully qualified domain name of the server for the CSR. Use a text string up to 64 characters in length.

KeyType

String

Indicates the type of key to generate. The following values are supported. The default value is RSA.

- RSA - generates RSA keys
- ECDSA - generates Elliptic Curve Digital Signature Algorithm (ECDSA) keys.

KeyLength

String

Indicates the length of the generated RSA keys. The following values are supported. The default value is 2048. The generation of a 4096-bit key can take up to 30 seconds. This keyword is required only when the key type set by KeyType is RSA.

- 1024 - generates a 1024-bit key.
- 2048 - generates a 2048-bit key.
- 4096 - generates a 4096-bit key.

Digest

String

Specifies the hash algorithm of the generated RSA keys. The default value is sha256. Digest is required only when KeyType is RSA and GenSSCert is set to TRUE.

- sha1 - uses SHA-1.
- sha256 - uses SHA-256.

ECDSACurve

String

Specifies the elliptic curve to use to generate the ECDSA keys. The curves in the following table are supported. The default value is secp256r1. This keyword is required only when KeyType is ECDSA.

<i>Table 67. Elliptic curves</i>		
Curve name	Value	Reference
sect163k1	1	RFC 4492

<i>Table 67. Elliptic curves (continued)</i>		
Curve name	Value	Reference
sect163k1	2	RFC 4492
sect163k1	3	RFC 4492
sect193r1	4	RFC 4492
sect193r1	5	RFC 4492
sect233k1	6	RFC 4492
sect233k1	7	RFC 4492
sect239k1	8	RFC 4492
sect283k1	9	RFC 4492
sect283k1	10	RFC 4492
sect409k1	11	RFC 4492
sect409k1	12	RFC 4492
sect571k1	13	RFC 4492
sect571k1	14	RFC 4492
secp160k1	15	RFC 4492
secp160k1	16	RFC 4492
secp160r2	17	RFC 4492
secp192k1	18	RFC 4492
secp192k1	19	RFC 4492
secp224k1	20	RFC 4492
secp224k1	21	RFC 4492
secp256k1	22	RFC 4492
secp256k1	23	RFC 4492
secp384r1	24	RFC 4492
secp521r1	25	RFC 4492
brainpoolP256r1	26	RFC 7027
brainpoolP384r1	27	RFC 7027
brainpoolP512r1	28	RFC 7027

FileName

String

Optionally specifies a common prefix for the generated private key, CSR, and self-signed certificate. If not specified, the value for the ObjectName parameter is used.

Days

String

Optionally specifies the validity period in days for the self-signed certificate. The default value is 365.

PasswordAlias

String

Optionally specifies the password alias map that defines the alias that maps to the cleartext password. The password in the map encrypts the files, and the alias in the map decrypts the passwords to access the file.

ExportKey

Enum

Optionally creates a copy of the private key in the `temporary: directory` in addition to the one in the `cert: directory`.

GenSSCert

Enum

Optionally creates a self-signed certificate in addition to the private key and CSR.

ExportSSCert

Enum

Optionally creates a copy of the self-signed certificate in the `temporary: directory` in addition to the one in the `cert: directory`.

GenObject

Enum

Creates a key management object. To create a certificate management object use the `GenSSCert` parameter.

ObjectName

Enum

Optionally specifies the names for the objects that are created by the `GenObject` parameter. If not specified, the value for the `CN` parameter is used.

HSM

Enum

Always set to "no".

UsingKey

String

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.

ConfigFile

String

Optionally, when generating a CSR or self-signed certificate, specifies the configuration file to read the DN and OIDs from. The configuration file must be in the OpenSSL CNF format. You must upload the configuration file to the `temporary: directory`.

When you specify the `ConfigFile` parameter, you must specify the `CN`, `FileName`, or `ObjectName` parameter.

Example

The following example generates 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 password alias `didgeridoo`.

To generate the key, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "Keygen": {
```

```

    "C": "au",
    "ST": "Victoria",
    "L": "South Melbourne",
    "O": "DataPower Australia, Ltd.",
    "OU": "Customer Support",
    "CN": "www.bob.datapower.com.au",
    "KeyType": "rsa",
    "KeyLength": "2048",
    "FileName": "Bob",
    "PasswordAlias": "didgeridoo",
  }
}

```

LoadConfiguration

Trigger this action to execute a number of other actions in one POST operation.

The request payload contains multiple objects and files that are otherwise used in single POST requests. The objects and files have the same format as they have if they are used in their dedicated actions.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```

{
  "LoadConfiguration": {
    "action_1": {
      action_parameters,
    }
    "action_2": {
      action_parameters,
    }
    ..
    "action_n": {
      action_parameters,
    }
  }
}

```

action

String

The action to include in the request.

action_parameters

Parameters required for the requested action.

Example

This LoadConfiguration payload uses three kinds of object and file information to create a file.

1. Set values to create the object by using the CryptoKey name.
2. Specify the file to create by using the file name.
3. Provide a password alias for the file to create.

```

{
  "LoadConfiguration": {
    "CryptoKey": {
      "name": "DEFAULTKEY",
      "mAdminState": "enabled",
      "Filename": "sharedcert:///default.p12",
      "PasswordAlias": "a4"
    },
    "file": {
      "path": "/sharedcert/default.p12",
      "content": "PCEtLSB0aGlzIGZpbGUgY29udGFpbmMgc3RhdG1jIHN0cm1uZ3MgZGVmaW5lZCBieSB0
aGUgV1MtQWRkcmVzc2luZyBzcGVjaWZpY2F0aW9uIC0tPgoKPCFFt1RjVfkgd3NhLWZhdWx0ICJodHRwOi8vc2
NoZW1hcy54bWxzbnZlLm9yZy93cy8yMDA0LzA4L2FkZHZJc3NpbmVzc2F1bHQiPgo8IUV0VE1UWSB3c2E0YW5v
biAgImh0dHA6Ly9yZ2h1bWZzLnhtbHNvYXAub3JnL3dzLzIwMDQvMDgvdWYWRkcmVzc2luZy9yb2x1L2Fub255bW
91cyI+CgoKCg=="
    },
    "PasswordAlias": {
      "name": "a4",
      "mAdminState": "enabled",
      "Password": "p1"
    }
  }
}

```

```
}  
}
```

LocateDevice

Trigger this action to light the blue LED on the front of the appliance to help you locate it in a data centre.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{  
  "LocateDevice": {  
    "LocateLED": "on" | "off"  
  }  
}
```

Example

To turn the blue LED on your appliance on, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{  
  "LocateDevice": {  
    "LocateLED": "on"  
  }  
}
```

MoveFile

Trigger this action to move or rename a file on the appliance.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{  
  "MoveFile": {  
    "sURL": "filepath",  
    "dURL": "filepath",  
    "Overwrite": "yes" | "no"  
  }  
}
```

sURL

String

Specifies the location of the source file as a URL in the *directory:///file* format.

dURL

String

Specifies the location of the destination file as a URL in the *directory:///file* format.

Overwrite

Enum

Set to "yes" to specify that the destination file can be overwritten if it already exists.

Examples

To move the file `config:///startup-config` to `store:///archiveConfig-10`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{  
  "MoveFile": {  
    "sURL": "config:///startup-config",  
    "dURL": "store:///archiveConfig-10",  
    "Overwrite": "yes"  
  }  
}
```

To rename the file `config:///startup-config` to `store:///archiveConfig-10`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "MoveFile": {
    "sURL": "config:///startup-config",
    "dURL": "config:///archiveConfig-10",
    "Overwrite": "yes"
  }
}
```

NoPasswordMap

Trigger this action to clear the entries in the password map.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "NoPasswordMap": {}
}
```

Ping

Trigger this action to ping a remote host.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "Ping": {
    "RemoteHost": "host",
    "useIPv": "IP_address_format"
  }
}
```

RemoteHost

String

Specify a host name or IP address.

useIPv

String

Optionally specify the format of the IP address. Specify "-4" for IPv4 (the default), or "-6" for IPv6.

The appliance sends Internet Control Message Protocol (ICMP) echo-request messages to the remote host to test connectivity.

Example

To ping the host `williamswillow.co.uk`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "Ping": {
    "RemoteHost": "williamswillow.co.uk",
  }
}
```

RaidActivate

Trigger this action to activate a RAID volume.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "RaidActivate": {
    "VolumeName": "name"
  }
}
```

```
}  
}
```

VolumeName

String

Specify the name of the RAID volume that you want to activate. The name is always `raid0`.

Example

To activate the `raid0` volume, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{  
  "RaidActivate": {  
    "VolumeName": "raid0"  
  }  
}
```

RaidDelete

Trigger this action to delete a RAID volume.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{  
  "RaidDelete": {  
    "VolumeName": "name"  
  }  
}
```

VolumeName

String

Specify the name of the RAID volume that you want to delete. The name is always `raid0`.

This action deletes all files and data on the RAID array.

Example

To delete the `raid0` volume, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{  
  "RaidDelete": {  
    "VolumeName": "raid0"  
  }  
}
```

RaidInitialize

Trigger this action to initialize a RAID volume.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{  
  "RaidInitialize": {  
    "VolumeName": "name"  
  }  
}
```

VolumeName

String

Specify the name of the RAID volume that you want to initialize. The name is always `raid0`.

This action deletes any files and data on the RAID array.

Example

To initialize the raid0 volume, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "RaidInitialize": {
    "VolumeName": "raid0"
  }
}
```

RaidLearnBattery

Trigger this action to request the BBU to start the learning cycle..

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "RaidLearnBattery": {
  }
}
```

This action requests the battery backup unit (BBU) to start the learning cycle. This action takes approximately six hours to complete. During the learning cycle, the write cache is disabled. During this period, write performance is slower.

RaidMakeHotSpare

Trigger this action to create a hot spare for a RAID volume.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "RaidMakeHotSpare": {
    "VolumeName": "name"
  }
}
```

VolumeName

String

Specify the name of the RAID volume that you want to make a hot spare for. The name is always `raid0`.

This action 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

To make a hot spare, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "RaidMakeHotSpare": {
    "VolumeName": "raid0"
  }
}
```

RaidRebuild

Trigger this action to force a rebuild of a RAID volume.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "RaidRebuild": {
```

```
    "VolumeName": "name"
  }
}
```

VolumeName

String

Specify the name of the RAID volume that you want to rebuild. The name is always `raid0`.

This action forces a rebuild of a RAID volume. The contents of the primary disk are copied to the secondary disk.

Example

To rebuild the RAID volume, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "RaidRebuild": {
    "VolumeName": "raid0"
  }
}
```

RemoveDir

Trigger this action to remove subdirectories from the `local`: directory.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "RemoveDir": {
    "Dir": "local:///subdirectory"
  }
}
```

Dir

String

The subdirectory to remove from the `local`: directory.

Example

To remove the , post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "RaidRebuild": {
    "VolumeName": "raid0"
  }
}
```

SaveConfig

Trigger this action to save a running configuration to memory.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "SaveConfig": {
  }
}
```

SaveInternalState

Trigger this action to save the internal state as a text file.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "SaveInternalState": {}
}
```

The action writes the internal state to the temporary `:///internal-state.txt` file.

SelectConfig

Trigger this action to designate the startup configuration for the next restart.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "SelectConfig": {
    "File": "file"
  }
}
```

File

String

Specifies the name of the startup configuration file.

Example

To designate `testEnvironment.cfg` as the startup configuration for the next restart, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "SelectConfig": {
    "File": "testEnvironment.cfg"
  }
}
```

SendErrorReport

Trigger this action to send an error report as email.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "SendErrorReport": {
    "SmtpServer": "server",
    "LocationIdentifier": "subject",
    "EmailAddress": "file",
    "EmailSenderAddress": "format"
  }
}
```

SmtpServer

String

Identifies a local SMTP server by IP address or by host name.

LocationIdentifier

String

Specifies the text string for the subject field of the outgoing message.

EmailAddress

String

Specifies the full email address of the message recipient.

EmailSenderAddress

String

Specifies the full email address of the message sender.

Example

To send an error report in an email to the address billgamsworth@paddock.com, post the following payload to https://yourhost:5554/mgmt/actionqueue/default:

```
{
  "SendErrorReport": {
    "SmtpServer": "smtp.paddock.com",
    "LocationIdentifier": "error report",
    "EmailAddress": "billgamsworth@paddock.com",
    "EmailSenderAddress": "JoeyMidtown@paddock.com"
  }
}
```

SendFile

Trigger this action to send a file using SMTP-based file transmission.

Use the HTTP POST method with the resource /mgmt/actionqueue/default, specifying a request payload with the following format:

```
{
  "SendFile": {
    "LocationIdentifier": "subject",
    "SmtpServer": "server",
    "EmailAddress": "file",
  }
}
```

LocationIdentifier

String

Identifies the target file and takes one of the following forms:

- audit:///file
- pubcert:///file
- config:///file
- store:///file
- image:///file
- tasktemplates:///file
- logstore:///file
- temporary:///file
- logtemp:///file

SmtpServer

String

Identifies a local SMTP server by IP address or by host name.

EmailAddress

String

Specifies the full email address of the message recipient.

Example

To send the file temporary:///stuff to the address billgamsworth@paddock.com, post the following payload to https://yourhost:5554/mgmt/actionqueue/default:

```
{
  "SendFile": {
```

```
"LocationIdentifier": "temporary:///stuff",
"SmtServer": "smtp.paddock.com",
"EmailAddress": "billgamsworth@paddock.com"
}
}
```

SendLogEvent

Trigger this action to verify that log targets are active and able to capture events.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "SendLogEvent": {
    "LogType": "category",
    "GenLogLevel": "priority",
    "LogEvent": "message",
    "EventCode": "code"
  }
}
```

LogType

String

Specifies the name of a log category.

GenLogLevel

String

Sets the event priority. The priority indicates that all events that are greater than or equal to this value are logged. Events use the following priority in descending order.

- For emergency, use `emerg` or 0.
- For alert, use `alert` or 1.
- For critical, use `critic` or 2.
- For error, use `error` or 3.
- For warning, use `warn` or 4.
- For notice, use `notice` or 5.
- For information, use `info` or 6.
- For debug, use `debug` or 7.

LogEvent

String

Specify the message text.

EventCode

String

Optionally specify an error code.

This action generates an event at a specific priority and sends it to the system log. Use this command to test the configuration of a log target.

SetLogLevel

Trigger this action to specify the minimum level of events that are captured in the system log.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "SetLogLevel": {
    "LogLevel": "level",
  }
}
```

LogLevel

String

Sets the event priority. The priority indicates that all events that are greater than or equal to this value are logged. Events use the following priority in descending order.

- For emergency, use `emerg` or 0.
- For alert, use `alert` or 1.
- For critical, use `critic` or 2.
- For error, use `error` or 3.
- For warning, use `warn` or 4.
- For notice, use `notice` or 5.
- For information, use `info` or 6.
- For debug, use `debug` or 7.

Example

To set the log level so that events of error and above are written to the log, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "SetLogLevel": {
    "LogLevel": "3",
  }
}
```

SetRBMDDebugLog

Trigger this action to enable RBM debug logging.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "SetRBMDDebugLog": {
    "RBMLog": "yes" | "no"
  }
}
```

RBMLog

Enum

Specify "yes" to enable RBM debug logging in the system log.

Example

To enable RBM logging, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "SetRBMDDebugLog": {
    "RBMLog": "yes"
  }
}
```

SetTimeAndDate

Trigger this action to set the time or the date on the appliance.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "SetTimeAndDate": {
    "Date": "date",
  }
}
```

```
    "Time": "time"
  }
}
```

Date

String

Specify the date in the format *yyyy-mm-dd*.

Time

String

Specify the time in the format *hh:mm:ss*.

Example

To set a time of 9:45 am and a date of the 18th August 2020, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "SetTimeAndDate": {
    "Date": "2020-08-18",
    "Time": "09:45:00"
  }
}
```

Shutdown

Trigger this action to shutdown or restart the appliance.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "Shutdown": {
    "Mode": "mode",
    "Delay": seconds
  }
}
```

Mode

String

Specify the shutdown mode:

- `reload` - Reloads the firmware.
- `reboot` - Restarts the appliance.
- `poweroff` - Stops the appliance and turns off the power.

Delay

Int

Specifies the number of seconds before a shutdown operation starts. Enter a value in the range 0 - 65535. The default value is 10.

Example

To stop the appliance and turn of the power in two minutes time, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "Shutdown": {
    "Mode": "poweroff",
    "Delay": 120
  }
}
```

TCPConnectionTest

Trigger this action to test connectivity to a remote host using a specific port.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "TCPConnectionTest": {
    "RemoteHost": "host",
    "RemotePort": "port",
    "Timeout": seconds,
    "useIPv": "IP_address_format",
    "LocalAddress": "IP_address"
  }
}
```

RemoteHost

String

Specify a host name or IP address.

RemotePort

String

Specify the port number.

Timeout

Int

Specify the timeout for this connection attempt in seconds.

useIPv

String

Optionally specify the format of the IP address. Specify "-4" for IPv4 (the default), or "-6" for IPv6.

LocalAddress

String

Optionally specify the IP address of a local interface to use for the connection test.

Example

To test the connection to port 5554 of the host `williamswillow.co.uk` with a twenty second timeout, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "TCPConnectionTest": {
    "RemoteHost": "williamswillow.co.uk",
    "RemotePort": "5554",
    "Timeout": 20
  }
}
```

TestHardware

Trigger this action to test the appliance hardware.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "TestHardware": {
  }
}
```

The components tested are grouped in the following categories:

- Backtrace availability
- Interface diagnostics
- Fan diagnostics

- RAID volume diagnostics
- Sensors diagnostics
- CPU/memory diagnostics

Results are returned in the following categories:

- success
- warning
- failure

TraceRoute

Trigger this action to trace the network path to a remote host.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "TraceRoute": {
    "RemoteHost": "host",
    "useIPv": "IP_address_format"
  }
}
```

RemoteHost

String

Specify a host name or IP address.

useIPv

String

Optionally specify the format of the IP address. Specify "-4" for IPv4 (the default), or "-6" for IPv6.

Example

To trace the path to the host `williamswillow.co.uk`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "TraceRoute": {
    "RemoteHost": "williamswillow.co.uk",
  }
}
```

UndoConfig

Trigger this action to revert to the last saved configuration.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "UndoConfig": {
    "Class": "config_type",
    "Name": "config_name"
  }
}
```

Class

String

The type of configuration that you are reverting.

Name

String

The name of the configuration that you are reverting.

Example

To , post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "UndoConfig": {
    "Class": "config_type",
    "Name": "config_name"
  }
}
```

UniversalPacketCaptureDebug

Trigger this action to manage a packet capture session. Packet capture is a tool that troubleshoots network connectivity, TCP sequencing, and other network-level problems.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "UniversalPacketCaptureDebug": {
    "Filename": "filepath",
    "InterfaceType": "interface_type",
    "EthernetInterface": "Eth_interface_name",
    "VLANInterface": "VLAN_interface_name",
    "LinkAggregation": "linkagg_interface_name",
    "StandaloneStandbyControlInterface": "standalone_interface_name",
    "CaptureMode": "mode",
    "MaxTime": seconds,
    "MaxSize": KB,
    "MaxPacketSize": size,
    "Filter": "expression",
    "SSLKeyLog": "yes" | "no"
  }
}
```

Filename

String

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.

InterfaceType

String

Specifies the type of interface the packet capture is for. The following types are available:

- **Ethernet** - Indicates that the capture is on the base Ethernet MAC interface. Traffic over any VLAN sub-interface configured on the Ethernet interface is included in the capture.
- **VLAN** - Indicates that the capture is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.
- **Loopback** - Indicates that the capture is for localhost traffic. This traffic is on the software loopback interface for Internet Protocol traffic (CIDR 127/8 in IPv4 or ::1/128 in IPv6).
- **Aggregate** - Indicates that the capture is on an IEEE 802.3ad link aggregation interface, which is composed of one of the base Ethernet MAC interface.
- **Standalone** - Indicates that the capture is on standalone interfaces which are configured by the underlying operating system.
- **All** - Indicates that the capture is on all Ethernet MAC interfaces. Packets that arrive on any Ethernet MAC interfaces are captured.

EthernetInterface

String

If you have specified an `InterfaceType` of `Ethernet`, specify the name of the Ethernet interface you are setting up packet capture for.

VLANInterface

String

If you have specified an `InterfaceType` of `VLAN`, specify the name of the VLAN interface you are setting up packet capture for.

LinkAggregation

String

If you have specified an `InterfaceType` of `Aggregate`, specify the name of the Link Aggregation interface you are setting up packet capture for.

StandaloneStandbyControlInterface

String

If you have specified an `InterfaceType` of `Standalone`, specify the name of the stand alone interface you are setting up packet capture for.

CaptureMode

String

Specify the timing mode for the packet capture:

- `timed` - The packet capture stops after the specified time or earlier if another criterion is met.
- `continuous` - The capture continues until it is manually stopped.

MaxTime

Int

Maximum duration of the packet capture in timed mode. The minimum value is 5 seconds and the maximum value is 86400 seconds.

MaxSize

Int

The maximum size of the entire packet capture. If this size is reached the packet capture stops. Enter a value in the range 10 - 500000.

MaxPacketSize

Int

The maximum size of each packet recorded during a packet capture session. Enter any value in the range 20 - 9000 bytes. The special value `-1` invokes the default 9000 bytes, which corresponds to the maximum transmission unit for Ethernet interfaces and therefore allows the entire frame to be saved.

Filter

String

The packet capture filter expression is a string that defines how the packet capture should be filtered. Full details of the syntax can be found by searching the web for "`pcap-filter(7)` format".

SSLKeyLog

Boolean

Indicates whether to capture the SSL master secret for SSL sessions. When you capture the SSL master secrets, the `logtemp:///sslkeyfile.log` file contains the master secrets. You can use a network protocol analyzer to decrypt SSL encrypted traffic to help with problem determination. During the packet capture, all SSL sessions have their master secret logged.

Example

The following example starts 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.

To start the example packet capture session on the VLAN interface `VLAN203`, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "UniversalPacketCaptureDebug": {
    "Filename": "temporary:///capture-2",
    "InterfaceType": "VLAN",
    "VLANInterface": "VLAN203",
    "CaptureMode": "timed",
```

```

    "Time": 1800,
    "Size": 2500,
    "Expression": "dst port 53",
  }
}

```

UniversalStopPacketCapture

Trigger this action to stop a packet capture session.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```

{
  "UniversalStopPacketCapture": {
    "Filename": "filepath",
    "InterfaceType": "interface_type",
    "EthernetInterface": "Eth_interface_name",
    "VLANInterface": "VLAN_interface_name",
    "LinkAggregation": "linkagg_interface_name",
    "StandaloneStandbyControlInterface": "standalone_interface_name"
  }
}

```

Filename

String

Specifies the file name for the packet capture that you want to stop.

InterfaceType

String

Specifies the type of interface the packet capture is for. The following types are available:

- **Ethernet** - Indicates that the capture is on the base Ethernet MAC interface. Traffic over any VLAN sub-interface configured on the Ethernet interface is included in the capture.
- **VLAN** - Indicates that the capture is on an IEEE 802.1Q VLAN on top of the base Ethernet MAC interface.
- **Loopback** - Indicates that the capture is for localhost traffic. This traffic is on the software loopback interface for Internet Protocol traffic (CIDR 127/8 in IPv4 or ::1/128 in IPv6).
- **Aggregate** - Indicates that the capture is on an IEEE 802.3ad link aggregation interface, which is composed of one of the base Ethernet MAC interface.
- **Standalone** - Indicates that the capture is on standalone interfaces which are configured by the underlying operating system.
- **All** - Indicates that the capture is on all Ethernet MAC interfaces. Packets that arrive on any Ethernet MAC interfaces are captured.

EthernetInterface

String

If you have specified an `InterfaceType` of `Ethernet`, specify the name of the Ethernet interface you are stopping the packet capture for.

VLANInterface

String

If you have specified an `InterfaceType` of `VLAN`, specify the name of the VLAN interface you are stopping the packet capture for.

LinkAggregation

String

If you have specified an `InterfaceType` of `Aggregate`, specify the name of the Link Aggregation interface you are stopping the packet capture for.

StandaloneStandbyControlInterface

String

If you have specified an `InterfaceType` of `Standalone`, specify the name of the stand alone interface you are stopping the packet capture for.

Example

The following example stops the packet-capture session that writes data to the `temporary:///capture-2` file on the VLAN interface named `VLAN203`.

To stop the packet capture session, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "UniversalStopPacketCapture": {
    "Filename": "temporary:///capture-2",
    "InterfaceType": "VLAN",
    "Interface": "VLAN203"
  }
}
```

UserForcePasswordChange

Trigger this action to force a user to change their password the next time they log in to the appliance.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "UserForcePasswordChange": {
    "User": "user_id"
  }
}
```

User

String

The user who must change their password.

Example

To force the user `appliance_support` to change their password, post the following payload to `https://yourhost:5554/mgmt/actionqueue/default`:

```
{
  "UserForcePasswordChange": {
    "User": "appliance_support"
  }
}
```

VerifyFirmware

Trigger this action to run the integrity verifier tool against the current firmware image.

Use the HTTP POST method with the resource `/mgmt/actionqueue/default`, specifying a request payload with the following format:

```
{
  "VerifyFirmware": {
  }
}
```

When you run the integrity verifier tool, it verifies the integrity of files that are part of the currently installed firmware. When files are different than the currently installed firmware image, messages are written to the console and also written as warnings to the system log.

IBM MQ resources

The URIs used to administer IBM MQ are listed in the IBM MQ documentation.

See [REST API resources](#) in the IBM MQ Knowledge Center.

Messages

Visit the messages section in the IBM MQ documentation for help with IBM MQ diagnostic messages.

http://www.ibm.com/support/knowledgecenter/SSFKSJ_9.1.0/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.

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.



Part Number:

(1P) P/N: